Cardinal MiniLab Plug-in Suite

User's Manual

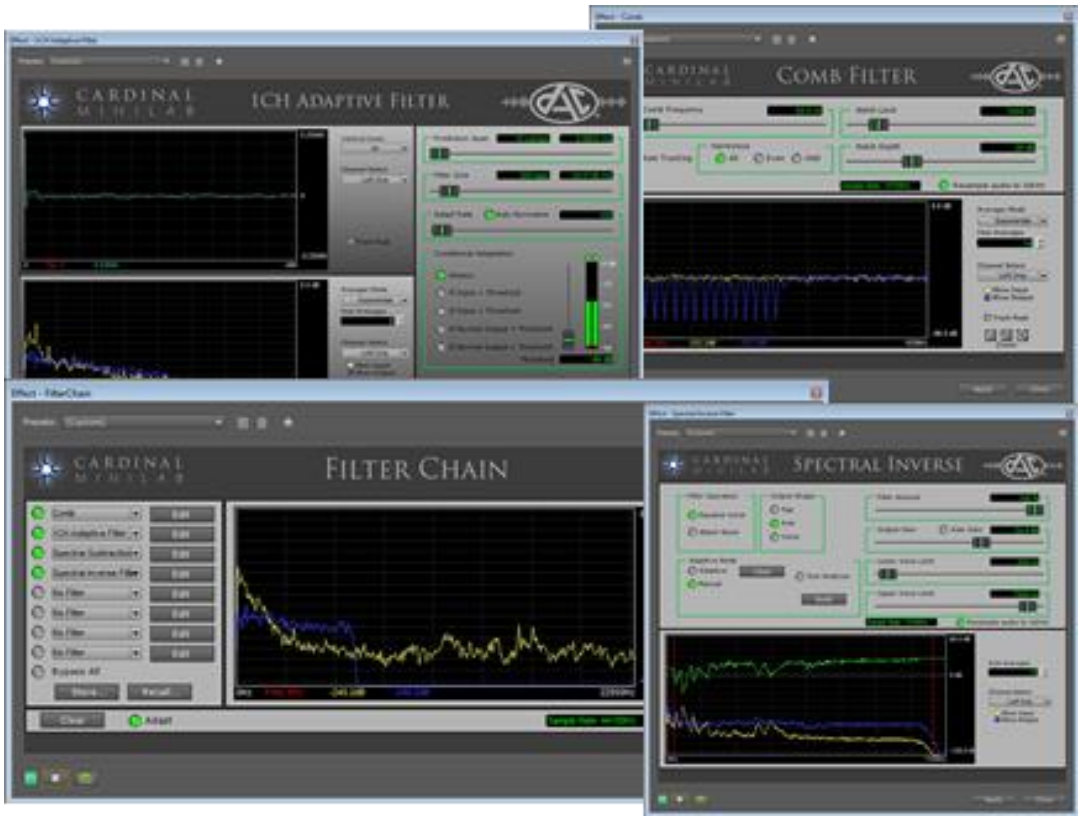

Salient Sciences

Clarity Revealed

Cardinal MiniLab Plug-in Suite

User's Manual

February 2015

Digital Audio Corporation
d/b/a Salient Sciences
4018 Patriot Drive Suite 300
Durham, NC 27703
Phone: +1 (919) 572 6767
Fax: +1 (919) 572 6786

sales@salientsciences.com
www.salientsciences.com

Copyright © 2015 by Salient Sciences.
All rights reserved.

TABLE OF CONTENTS

1 Cardinal MiniLab Introduction 9

2 Software Installation 10

2.1 Software Installation 10

2.2 Licensing And Activation 10

3 General Plug-in Concepts 14

3.1 Launching a Plug-in 14

3.1 Spectrum Analyzer 14

3.2 Coefficient display 16

3.3 Sample Rate Conversion 17

4 1-CH Adaptive Filter 18

4.1 Overview of the 1-CH Adaptive Filter 18

4.2 Main Plug-In Window 19

4.3 Filter Controls 19

4.3.1 Prediction Span 19

4.3.2 Filter Size 20

4.3.3 Adapt Rate 20

4.3.4 Conditional Adaptation 21

4.3.5 Filter Output 21

4.3.6 Freeze Filter 22

4.3.7 Clear Filter 22

5 Reference Canceller Filter 23

5.1 Overview of the Reference Canceller Filter 23

5.2 Main Plug-In Window 24

5.3 Filter Controls 24

5.3.1 Delay 24

5.3.2 Filter Size 25

5.3.3 Adapt Rate 25

5.3.4 Conditional Adaptation 26

5.3.5 Reference Gain 26

5.3.6 Filter Output 27

5.3.7 Freeze Filter 27

5.3.8 Clear Filter 27

6 Spectral Inverse Filter 28

6.1 Overview of the Spectral Inverse Filter 28

6.2 Main Plug-In Window 30

6.3 Filter Controls 31

6.3.1 Filter Operation 31
6.3.2 Output Shape 32
6.3.3 Filter Amount 32
6.3.4 Output Gain and Auto Gain 33
6.3.5 Lower and Upper Voice Limits 34
6.3.6 Adaptive Mode 35
6.3.7 Manual Mode 36
6.3.8 Display 36

7 Comb Filter 38

7.1 Overview of the Comb Filter 38

7.1 Main Plug-In Window 40

7.2 Filter Controls 41
7.2.1 Comb Frequency 41
7.2.2 Notch Limit 41
7.2.3 Notch Depth 41
7.2.4 Notch Harmonics 42
7.2.5 Auto Tracking 42

8 X-Pass Filter 43

8.1 Overview of the X-pass Filter 43

8.2 Main Plug-in Window 44

8.3 Lowpass Filter 45
8.3.1 Filter Controls 46

8.4 Highpass Filter 47
8.4.1 Filter Controls 48

8.5 Bandpass Filter 49
8.5.1 Filter Controls 50

8.6 Bandstop Filter 52
8.6.1 Filter Controls 53

8.7 Notch Filter 55
8.7.1 Filter Controls 56

8.8 Slot Filter 57
8.8.1 Filter Controls 58

9 Spectral Subtraction Filter 60

9.1 Overview of the Spectral Subtraction Filter 60

9.2 Main Plug-in Window 62

9.3 Filter Controls 63
9.3.1 Filter Controls 63
9.3.2 Clear 63
9.3.3 Filter Output Gain 64
9.3.4 Noise EQ Controls 64
9.3.5 Noise Reducer Controls 66

10 Multi-Band Filter 67
10.1 Overview of the Multi-Band Filter 67
10.1.1 Overview of Multiple-Notch 67
10.1.2 Overview of Multiple-Slot 68
10.2 Main Plug-in Window 70
10.3 Filter Controls 71
10.3.1 Filter Mode 71
10.3.2 Multiple Notch Filter Controls 71
10.3.3 Multiple Slot Controls 75
11 Graphic EQ 80
11.1 Overview of the Graphic Equalizer 80
11.1.1 Overview of Graphic EQ 80
11.1.2 Overview of Hi-Res Graphic 80
11.1 Main Plug-in Window 82
11.2 Equalizer Controls 83
11.2.1 Equalizer Mode 83
11.2.2 Graphic EQ Controls 83
11.2.3 Hi-Res Graphic Controls 84
12 Gain Stage 87
12.1 Overview of the Gain Stage 87
12.1.1 Overview of AGC 87
12.1.2 Overview of LCE 87
12.2 Main Plug-in Window 90
12.3 Gain Controls 91
12.3.1 Gain Mode 91
12.3.2 AGC Controls 91
12.3.3 LCE Controls 92
13 Parametric Equalizer 95
13.1 Overview of the Parametric Equalizer 95
13.2 Main Plug-in Window 96
13.3 Filter Controls 97
13.3.1 Current Stage 97
13.3.2 Center Frequency 97
13.3.3 Width Factor 97
13.3.4 Boost/Cut 98
13.3.5 Stage Active 98
13.3.6 Output Gain 98
13.3.7 Reset All 98
13.3.8 Show Inactive 99
13.3.9 Store Button 99
13.3.10 Recall Button 99

14 Filter Chain 99
14.1 Overview of the Filter Chain 99
14.2 Main Plug-in Window 100
14.3 Filter Chain Controls 101
14.3.1 Filter Select 101
14.3.2 Filter Control Button Group 101
14.3.3 Bypass All 102
14.3.4 Store Button 102
14.3.5 Recall Button 102
14.3.6 Clear Button 102

1 CARDINAL MINILAB INTRODUCTION

Thank you for purchasing the Cardinal MiniLab Plug-ins. The entire suite of plug-ins includes the following:

- 1-CH Adaptive Filter
- Comb Filter
- Filter Chain
- Gain Stage
- Graphic Equalizer
- Multi-Band Filter
- Parametric Equalizer
- Reference Canceller Filter
- Spectral Inverse Filter
- Spectral Subtraction Filter
- X-Pass Filter (6 filter set)

The MiniLab Suite includes all the above filters with the exception of the Reference Canceller and Filter Chain plug-ins. The Advanced Suite contains all the MiniLab plug-ins, plus an additional stand-alone application called Cardinal Batch Utility.

2 SOFTWARE INSTALLATION

2.1 SOFTWARE INSTALLATION

The software requires that a host audio editing environment that supports the VST standard be installed. The computer on which the audio editing environment is installed must meet all the requirements of that audio editing environment for the plug-in to work. Please consult the documentation of your audio editing environment for its requirements.

NOTE: Because of the processing requirements of the plug-in, the filters will benefit from a faster computer.

To start the installation, insert the CD containing the plug-in software into your PC. If the installation does not begin automatically run the file "Setup.exe" located in the root folder of the CD. You can run this file by clicking on the "Start" menu, then on "Run", and typing "X:\Setup.exe" (where "X" is the drive letter of your CD-ROM drive). Follow the instructions on the screen to complete the installation. If you have any questions or problems with the installation, please contact DAC for further assistance.

In addition to its own installation directory, the VST plug-in will also be copied to the default Steinberg VST plug-in folder (normally, "C:\Program Files\Steinberg\VstPlugins"). If this folder is not found, it will be created. Most audio editors will look in this directory by default for new plug-ins.

2.2 LICENSING AND ACTIVATION

When you first launch your plug-in, it will be running in Demo mode. This is normal operation until you enter the Serial Key you received from Digital Audio Corporation when you purchased your plug-in.

NOTE: In Demo mode, the plug-in will insert 1 second of silence for every 10 seconds of audio processed. There is nothing wrong with your plug-in, you simply need to enter your Serial Key to unlock the plug-in from Demo mode.

The plug-in itself provides a field to enter the Serial Key as shown in the following figure:

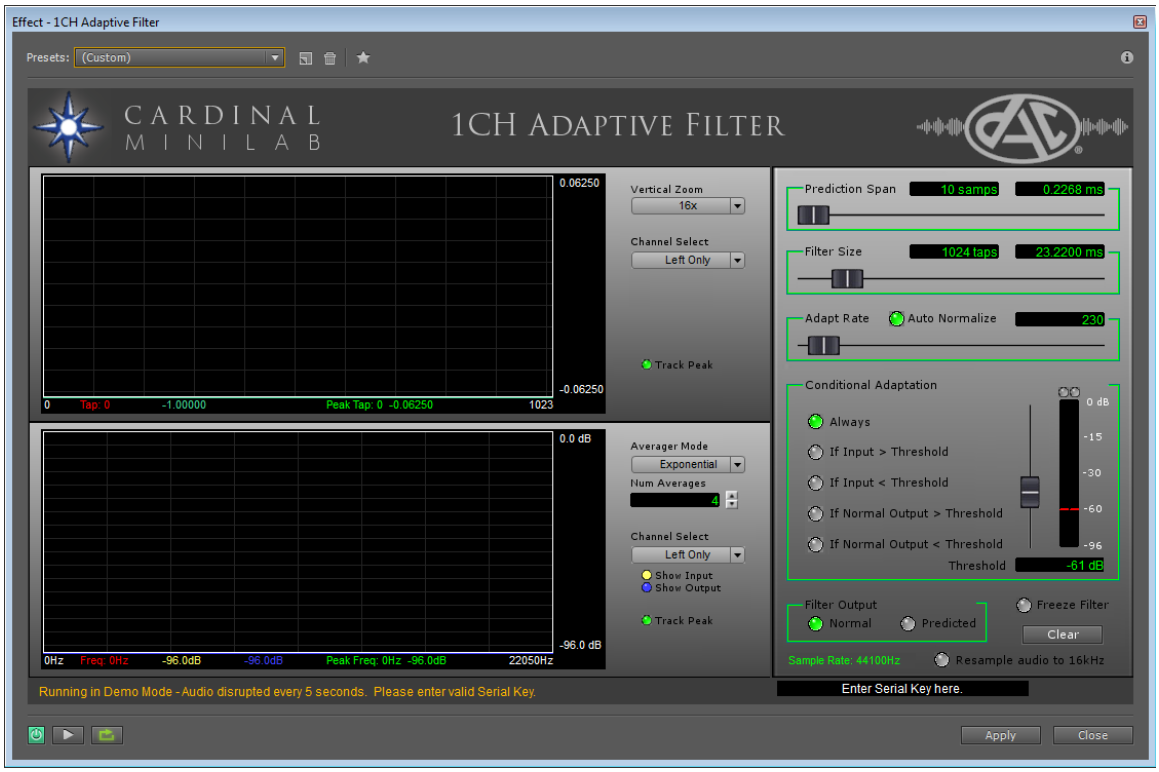

Figure 1: Plug-in in demo mode

Simply enter your Serial Key in the black text field provided in the lower-right of the screen and press Enter. The Serial Key MUST be entered exactly as provided by DAC. If the Serial Key is valid, it will be registered with your system and you will now have 30 days to activate the plug-in with DAC.

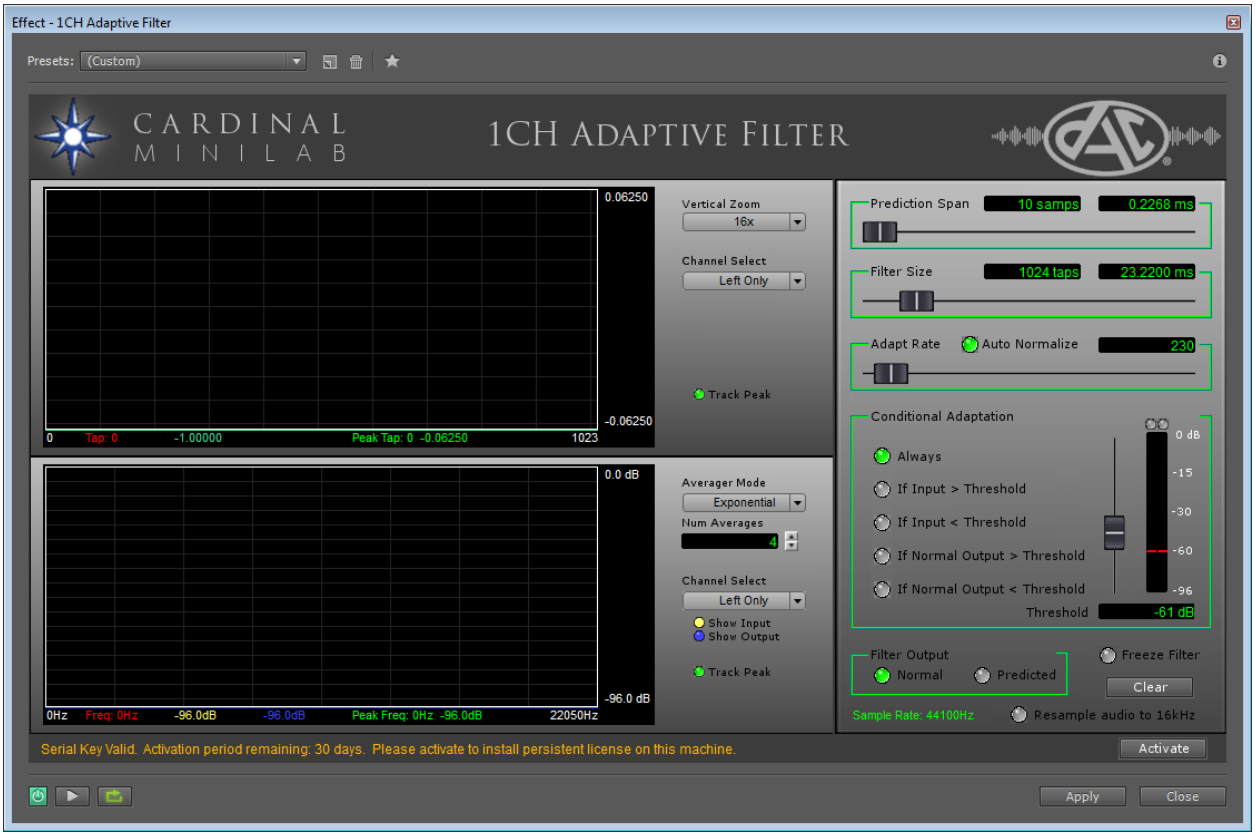

Figure 2: Plug-in waiting for Activation

To **Activate**, simply click the **Activate** button. This will send the necessary information to the DAC servers to activate the software on your PC. The ONLY information sent and stored at DAC is the Serial Key and a hardware identifier (usually a serial number from your System hard drive). ***No other personally identifiable information is sent.***

If the workstation you installed the software on does not have a network connection, the plug-in will display the Serial Key and HWID (Hardware Identifier) to you. You must then contact DAC with this information and we will generate a valid License Key for you that you can then enter as shown in the figure below.

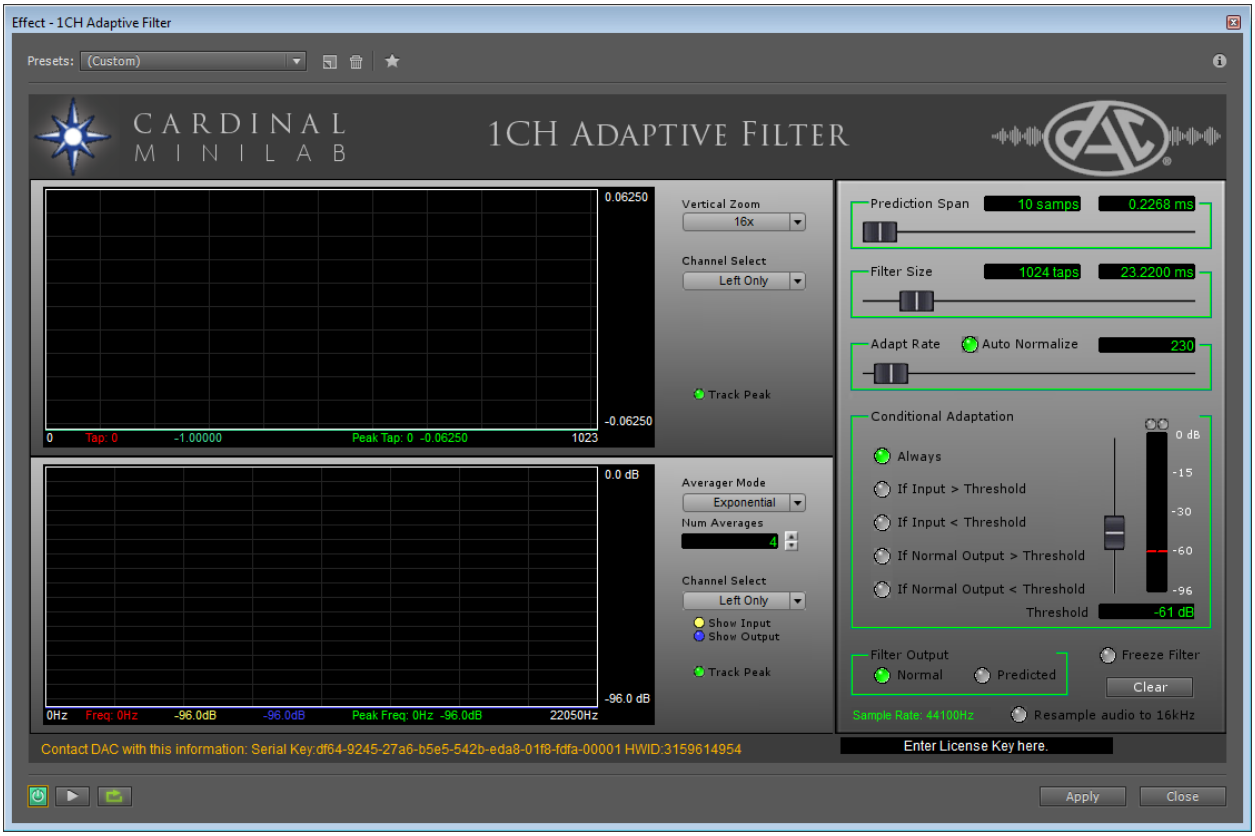

Figure 3: Plug-in waiting for manual License Key entry

Once activation is completed, either over the network or manually, a persistent license will be installed on the machine. If the software is not activated within 30 days, it will revert back to Demo mode.

You must uninstall the software (which will deactivate it as well) before installing the software on another machine.

If you have purchased one of the suites, you will only need to enter your serial key one time and activate it. All the other plug-ins in the suite will recognize the license and run normally. If you purchase plug-ins individually, then you will have to enter a separate serial key for each plug-in you purchased.

3 GENERAL PLUG-IN CONCEPTS

3.1 LAUNCHING A PLUG-IN

The Cardinal MiniLab suite of plug-ins are not stand-alone applications. They are plug-ins that execute from within a host environment such as Acon Digital's Acoustica, Adobe's Audition, Sony's SoundForge and other similar audio editors.

Plug-ins are launched from host applications in various ways depending on the host. Usually the plug-in can be found in the "Effects", "Tools" or "VST" menu. If the currently loaded audio file is not compatible with a plug-in, the plug-in choice will usually be grayed out (disabled).

In general, upon startup the host application will scan a well-known VST plug-in folder for valid VST plug-ins. Usually, new folders can be added to the scanned list.

Please consult your audio editor manual to determine how to recognize and launch a VST plug-in.

3.1 SPECTRUM ANALYZER

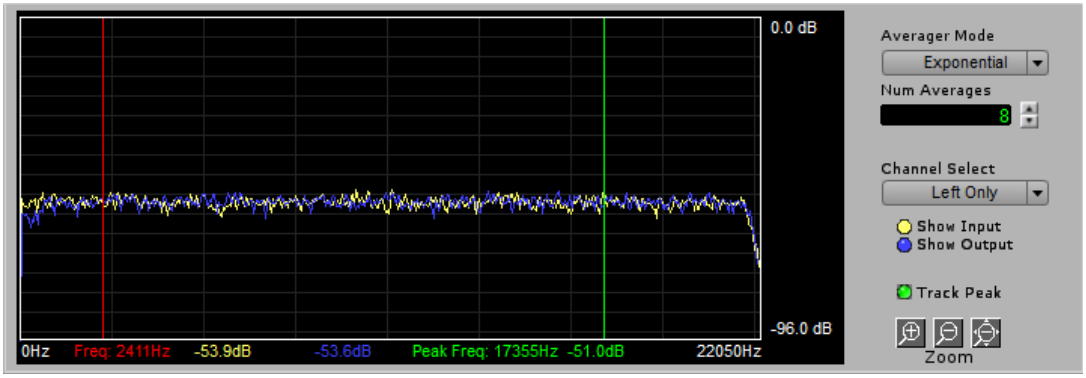

Figure 4: Spectrum Analyzer

Most plug-ins have an integrated spectrum analyzer to display frequency content of the audio. Most will look similar to the figure above.

The **Spectrum Analyzer** displays the frequency content of the input and/or output audio. The frequency axis of the **Spectrum Analyzer** plot goes from left to right, with the lowest frequency at the leftmost side and the highest frequency at the rightmost side. The "energy" (loudness) axis of the **Spectrum Analyzer** plot goes from top to bottom, with stronger (louder) frequencies

indicated by higher peaks on the plot.

For example, a signal consisting of a single tone will appear as a single peak, located at the tone frequency. A "white noise" signal, which contains equal amounts of every frequency in the signal bandwidth, will appear as an (approximately) flat line across the entire frequency range.

The spectrum analyzer can display the **Input** audio, the **Output** audio, or both. To enable the **Input** and/or **Output** trace, click the corresponding LED button so that the indicator "light" is on. The **Input** audio is the signal before the filter is applied. The **Output** audio is the signal after the filter is applied. When both signals are shown, each signal is indicated by a line, using yellow for the **Input** and blue for the **Output**.

The **Spectrum Analyzer** provides a cursor marker to help in identifying specific frequency values. At any time when the **Spectrum Analyzer** is activated, clicking in the graph area displays a red vertical marker at the frequency location clicked. The frequency (in Hz) of the location clicked is displayed in red text on the bottom left of the graph. The input level (in dB) is given in yellow text and the output level (in dB) is given in blue text.

The **Track Peak** feature displays a green vertical line in the graph display at the strongest (loudest) frequency. When this feature is enabled, the maximum peak value (in dB) and frequency (in Hz) is displayed in green text on the bottom right of the graph. To enable this indicator, click the button beside the **Track Peak** text so that the indicator “light” turns green.

Use the **Averager Mode** control to select what type of averaging is done to the signals displayed. **No Average** applies no averaging and the traces will move very rapidly. **Exponential** applies averaging to the trace values based on the **Num Averages** value. **Peak Hold** will hold the maximum value for each data point of the trace. Trace data will never leak down in this mode. This is useful for determining peak energy for a given frequency.

Use the **Num Averages** control to increase or decrease the number of averages applied to the frequency spectrum. A small **Num Averages** value allows a more accurate snapshot of the spectrum at a given time, but the trace values will change rapidly, making longer-term spectral characteristics difficult to see. Larger **Num Averages** values result in a smoother spectral plot that represents stable frequency characteristics well but does not accurately show rapidly time-varying signal characteristics. This control only has an effect when **Exponential** is selected for the **Averager Mode**.

The **Channel Select** control allows you to choose which channel of audio you wish to see displayed in the spectrum analyzer window. You can choose **Left**

Only, Right Only or **Mix**. The **Mix** selection sums the two audio channels together before performing the FFT analysis. **Right Only** and **Mix** are only available for stereo audio files.

The analyzer also has three **Zoom** controls: **Zoom In, Zoom Out** and **Zoom Full**. When zooming, the graph will always try to center the graph on wherever the red user marker is located. Each level of zoom halves the current frequency span of the display. For instance, if the current sample rate is 44100Hz, the first level of zoom will cause the analyzer to display 11025Hz instead of 22050Hz worth of spectral data. **Zoom Full** will always cause the analyzer to reset back to the full bandwidth of the signal.

3.2 COEFFICIENT DISPLAY

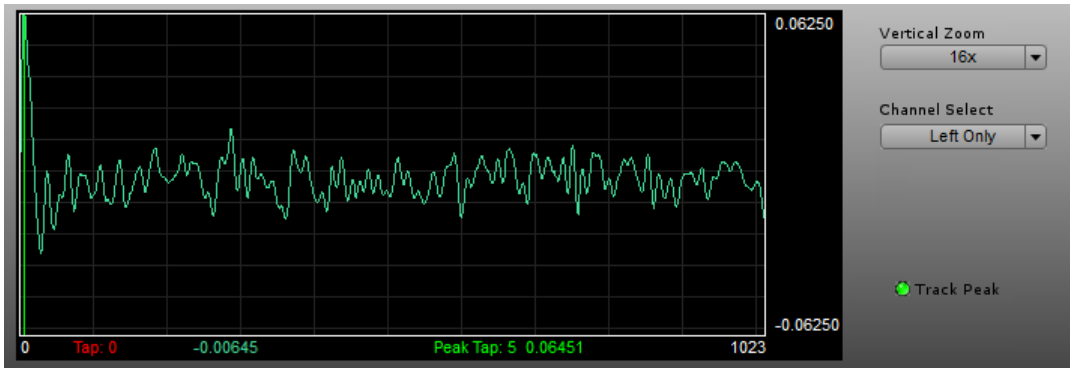

Figure 5: Coefficient Display

A few filters contain a Coefficient Display, particularly the 1-CH Adaptive Filter and Reference Canceller Filter. This graph can be useful in setting up these filters.

The Coefficient Display displays the impulse response (filter coefficients) of the filter. Vertical scaling of the filter's coefficients for display is accomplished by clicking on the **Vertical Zoom** combo box control. Supported zoom factors range from **1x** to **32x**.

The **Track Peak** feature displays a green vertical line in the graph display at the coefficient with the largest (absolute) value. When this feature is enabled, the maximum peak value and coefficient number is displayed in green text on the bottom right of the graph. To enable this indicator, click the button beside the **Track Peak** text so that the indicator "light" turns green.

The **Channel Select** control allows you to choose which channel's filter coefficients you wish to see displayed in the coefficient display window. You can choose **Left Only,** or **Right Only**. **Right Only** is only available for stereo audio files.

3.3 SAMPLE RATE CONVERSION

Figure 6: Sample Rate Conversion

All Cardinal MiniLab plug-ins contain a resampling feature that utilizes a true Whittaker-Shannon Interpolation to resample the incoming audio at 16kHz. This method eliminates any distortion components commonly found in cruder resampling techniques (such as linear interpolation).

Resampling the audio down at 16kHz sample rate primarily does two things for forensic audio:

1) Bandlimits the audio to the forensic voice spectrum (200 – 5000Hz)
2) Allows for larger filter sizes with less computational requirements. For example, a 1024-point filter at 48kHz will require roughly 3x the CPU resources that the same 1024-point filter will require at 16kHz.

Note: After processing, the plug-in resamples the audio back to the original sample rate to return it to the audio editor environment. If you process a 44.1kHz WAV file, you will still have a 44.1kHz WAV file in the end, but any information above 8kHz will have been eliminated.

4 1-CH ADAPTIVE FILTER

4.1 OVERVIEW OF THE 1-CH ADAPTIVE FILTER

The 1-Channel Adaptive filter is used to automatically cancel predictable and convolutional noises from the input audio. Predictable noises include tones, hum, buzz, engine/motor noise, and, to some degree, music. Convolutional noises include echoes, reverberations, and room acoustics.

The 1-CH Adaptive Filter is a forensic plug-in that has been specifically designed to work with VST host audio editing systems.

The 1-CH Adaptive Filter comes equipped with the following features:

- Filter sizes up to 8192 taps
- Prediction Span up to 32768 samples
- Auto-Normalizing adapt rate
- Conditional Adaptation
- Selectable Filter Output
- Coefficient Display (with max peak indicator)
- Dual-trace spectral analysis (with max peak indicator and averaging)

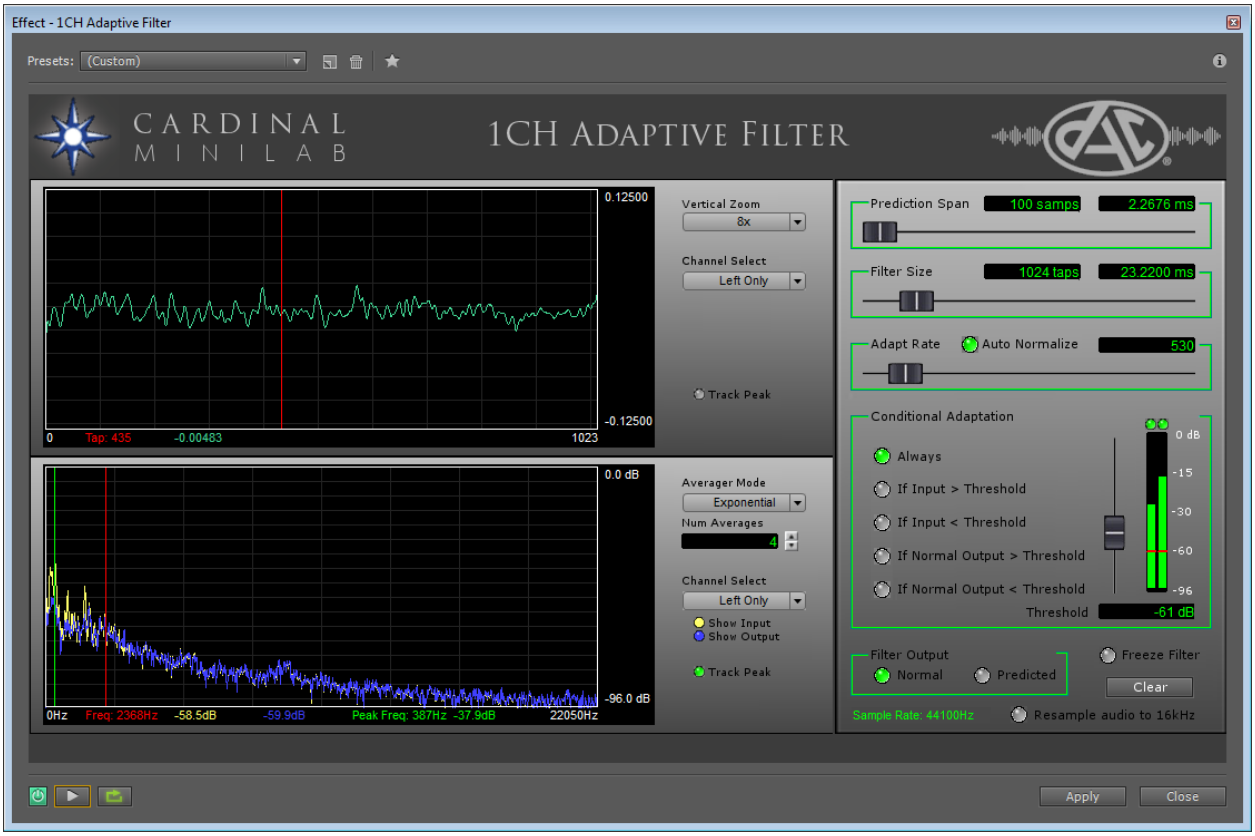

Figure 7: 1-CH Adaptive Filter Main Window

4.2 MAIN PLUG-IN WINDOW

Figure 7 shows the plug-in as it would be displayed if used in Adobe Audition. The buttons above and below the plug-in window will appear differently depending upon your audio editing environment. Please consult the documentation on your specific audio editing environment for details of how to use these features.

This window provides access to all the functionality of the plug-in. A coefficient display and spectrum analyzer are also provided as aides in determining the characteristics of the target audio. Details of the plug-in are described in the following sections.

4.3 FILTER CONTROLS

4.3.1 Prediction Span

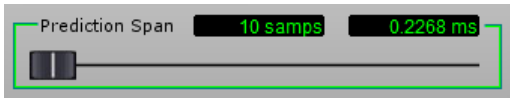

Figure 8: Prediciton Span

Sets the number of samples in the prediction span delay line. Prediction span is indicated both in samples and in milliseconds. Shorter prediction spans allow maximum noise removal, while longer prediction spans preserve voice naturalness and quality. *A prediction span of 2 or 3 samples is normally recommended.*

4.3.2 Filter Size

Figure 9: Filter Size

Sets the number of FIR filter taps in the adaptive filter. Filter size is indicated both in taps (filter order) and in milliseconds. The maximum filter size is 8192 taps. Small filters are most effective with simple noises such as tones and music. Larger filters should be used with complex noises such as severe reverberations and raspy power hums. *A nominal filter size of 512 to 1024 taps is a good overall general recommendation.*

CAUTION: Large filter sizes (> 2048 taps) will require large computing resources to maintain real-time audio processing. You may begin to hear skips in the audio during preview if your computer cannot keep up with the processing requirements. However, during render the audio will not contain any skips but may take longer than real-time to process the file.

4.3.3 Adapt Rate

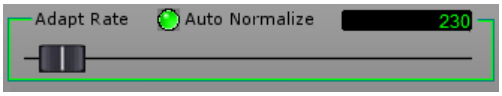

Figure 10: Adapt Rate

Used to set the rate at which the adaptive filter adapts to changing signal conditions. An adapt rate of 1 provides very slow adaptation, while an adapt rate of 5884 provides fastest adaptation. *A good approach is to start with an adapt rate of approximately 100-200 to establish convergence, and then back off to a smaller value to maintain cancellation.*
Larger adapt rates should be used with changing noises such as music; whereas, smaller adapt rates are acceptable for stable tones and reverberations. Larger adapt rates sometimes affect voice quality, as the filter may attack sustained vowel sounds.

When Auto Normalize is turned on, the specified adapt rate is continuously scaled based upon the input signal level. This scaling generally results in faster

filter convergence without greatly increasing the risk of a filter crash. *It is recommended that Auto Normalize be enabled for most speech signal processing.*

4.3.4 Conditional Adaptation

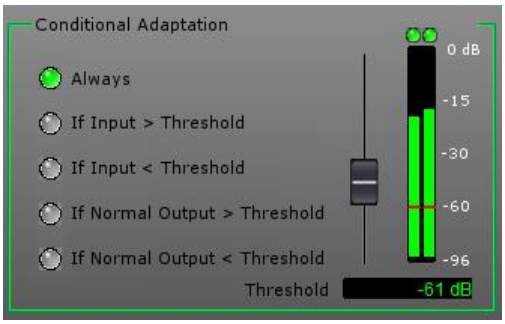

Figure 11: Conditional Adaptation

*For advanced users only. Novice users should keep Conditional Adaptation set to **Always**. The threshold setting has no effect when **Always** is selected.*

Conditional Adaptation allows the adaptive filter to automatically Adapt/Freeze based upon signal bargraph levels. This can be very useful in situations where there are pauses or breaks in the speech being processed.

Hint: Conditional adaptation is useful for maintaining adaptation once the filter has converged. Recording environment factors such as air temperature and motion in the room can cause the signal characteristics to change over the course of a recording. For this reason, simply freezing the filter once convergence is reached may mean that noise cancellation will degrade over time. Instead of freezing the filter, use Conditional Adaptation. First allow the filter to converge in **Always** mode, and then select **If Normal Output < Threshold**, and adjust the threshold by observing the bargraph levels during pauses in speech

Click on the **Clear** button if you desire the filter to completely readapt based upon the new Conditional Adaptation settings.

4.3.5 Filter Output

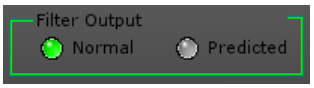

Figure 12: Filter Output

Used to optionally listen to the "rejected" audio that is being cancelled by the adaptive filter. ***Normal*** *should almost always be selected*, but the **Predicted** setting can be useful when configuring the filter, allowing the user to hear

exactly what is being removed by the filter.

4.3.6 Freeze Filter

Used to enable or disable filter adaptation. When Freeze is off, the filter adapts according to its settings. When Freeze is on, the filter never adapts regardless of the other settings.

4.3.7 Clear Filter

Used to reset the coefficients of the 1-CH Adaptive Filter. Clearing a filter is useful when the audio characteristics change dramatically, so that the filter can readapt to a new, clean solution. Clearing is also useful in the case of a filter "crash," when the filter coefficients diverge to an unstable state, usually in response to a large and abrupt change in the signal coupled with a fast adapt rate.

5 REFERENCE CANCELLER FILTER

5.1 OVERVIEW OF THE REFERENCE CANCELLER FILTER

The Reference Canceller adaptive filter is used to automatically cancel from the Primary channel any audio which matches the Reference channel. For example, the Primary channel may be microphone audio with desired voices masked by radio or TV noise. The radio/TV interference can be cancelled in real-time if the original broadcast audio, usually available from a second receiver, is simultaneously recorded to the Reference channel.

The Reference Canceller Filter is a forensic plug-in that has been specifically designed to work with VST host audio editing systems. Since it is an inherently two-channel filter, it is only available for stereo audio files. The Primary channel **MUST** be on the left channel, while the Reference channel consequentially **MUST** be on the right channel. Future updates of this plug-in will allow these channels to be switched if desired.

The Reference Canceller Filter comes equipped with the following features:

- Filter sizes up to 8192 taps
- Delay line up to 32768 samples
- Auto-Normalizing adapt rate
- Conditional Adaptation
- Reference Gain
- Selectable Filter Output
- Coefficient Display (with max peak indicator)
- Dual-trace spectral analysis (with max peak indicator and averaging)

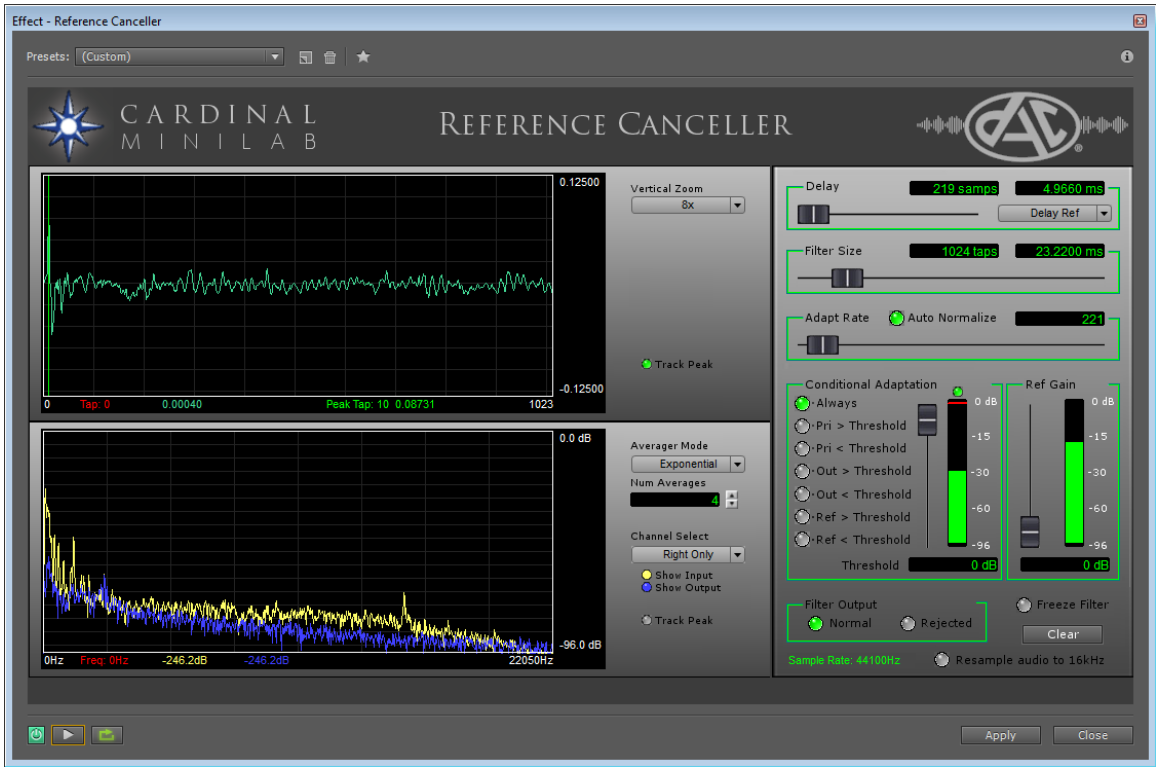

Figure 13: Reference Canceller Filter Main Window

5.2 MAIN PLUG-IN WINDOW

Figure 13 shows the plug-in as it would be displayed if used in Adobe Audition. The buttons above and below the plug-in window will appear differently depending upon your audio editing environment. Please consult the documentation on your specific audio editing environment for details of how to use these features.

This window provides access to all the functionality of the plug-in. A coefficient display and spectrum analyzer are also provided as aides in determining the characteristics of the target audio. Details of the plug-in are described in the following sections.

5.3 FILTER CONTROLS

5.3.1 Delay

Figure 14: Delay

Sets the number of audio samples by which the selected channel should be delayed. Adjusting the Delay allows the alignment of the Primary and Reference channels to be adjusted. Minimum Delay is 1 sample, but can be set to as high as 32768 samples.

Specifies whether the delay line is to go into either the Primary channel (**Delay Pri)** or the Reference channel (**Delay Ref**). *For most applications, a slight delay (typically 5 msec) is placed in the Primary channel*, For applications with long distances between the microphone and radio/TV, a delay in the Reference channel may be required. Extreme caution should be exercised when using reference channel delay; allowing the reference to lag the target noise in the primary signal will result in poor cancellation.

5.3.2 Filter Size

Figure 15: Filter Size

Sets the number of filter taps in the adaptive filter. Filter size is indicated both in taps (filter order) and in milliseconds. The maximum filter size is 8192 taps. *Normally, larger filters sizes are used in the Reference Canceller adaptive filter.*

CAUTION: Large filter sizes (> 2048 taps) will require large computing resources to maintain real-time audio processing. You may begin to hear skips in the audio during preview if your computer cannot keep up with the processing requirements. However, during render the audio will not contain any skips but may take longer than real-time to process the file.

5.3.3 Adapt Rate

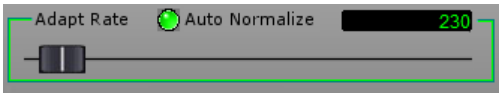

Figure 16: Adapt Rate

Used to set the rate at which the adaptive filter adapts to changing signal conditions. An adapt rate of 1 provides very slow adaptation, while an adapt rate of 5884 provides fastest adaptation. *A good approach is to start with an adapt rate of approximately 100-200 to establish convergence, and then back off to a smaller value to maintain cancellation.*

When Auto Normalize is turned on, the specified adapt rate is continuously scaled based upon the input signal level. This scaling generally results in faster

filter convergence without greatly increasing the risk of a filter crash. *It is recommended that Auto Normalize be enabled for most speech signal processing.*

5.3.4 Conditional Adaptation

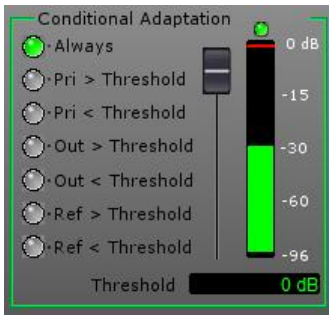

Figure 17: Conditional Adaptation

*For advanced users only. Novice users should keep Conditional Adaptation set to **Always**. The threshold setting has no effect when **Always** is selected.*

Conditional Adaptation allows the adaptive filter to automatically Adapt/Freeze based upon signal bargraph levels. This can be very useful in situations where there are pauses or breaks in the speech being processed.

Hint: Conditional adaptation is useful for maintaining adaptation once the filter has converged. Recording environment factors such as air temperature and motion in the room can cause the signal characteristics to change over the course of a recording. For this reason, simply freezing the filter once convergence is reached may mean that noise cancellation will degrade over time. Instead of freezing the filter, use Conditional Adaptation. First allow the filter to converge in **Always** mode, and then select **If Normal Output < Threshold**, and adjust the **Threshold** by observing the bargraph levels during pauses in speech

Click on the **Clear** button if you desire the filter to completely readapt based upon the new Conditional Adaptation settings.

5.3.5 Reference Gain

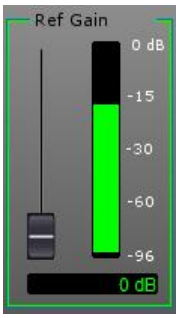

Figure 18: Reference Gain

Ref Gain is used to add gain to the reference channel audio if necessary. *To achieve good cancellation, it is important that the reference audio be at least as loud as the noise it is intended to cancel from the primary audio.*

5.3.6 Filter Output

Figure 19: Filter Output

Used to optionally listen to the "rejected" audio that is being cancelled by the adaptive filter. *Normal should almost always be selected,* but the Rejected setting can be useful when configuring the filter, allowing the user to hear exactly what is being removed by the filter.

5.3.7 Freeze Filter

Used to enable or disable filter adaptation. When Freeze is off, the filter adapts according to its settings. When Freeze is on, the filter never adapts regardless of the other settings.

5.3.8 Clear Filter

Used to reset the coefficients of the Reference Canceller Filter. Clearing a filter is useful when the audio characteristics change dramatically, so that the filter can readapt to a new, clean solution. Clearing is also useful in the case of a filter "crash," when the filter coefficients diverge to an unstable state, usually in response to a large and abrupt change in the signal coupled with a fast adapt rate.

6 SPECTRAL INVERSE FILTER

6.1 OVERVIEW OF THE SPECTRAL INVERSE FILTER

The Spectral Inverse Filter (SIF) has two modes: an adaptive mode and a manual mode. Thus it is actually a combination of the Adaptive Spectral Inverse Filter (ASIF) and the traditional Spectral Inverse Filter (SIF). Both are equalization filters that readjust the spectrum to match an expected spectral shape. It is especially useful when the target voice has been exposed to spectral coloration (i.e. muffling, hollowness, or tinniness), but it can also be used to remove bandlimited noises. The ASIF is much like the SIF, except it continually updates the spectral solution, whereas the SIF only updates the solution when it is "built".

The filter maintains an average of the signal's spectrum and uses this information to implement a high-resolution digital filter for correcting long-term spectral irregularities. The goal of the filter is to reshape the overall spectral envelope of the audio, not to respond to transient noises and characteristics.

Several user controls are available for refinement of SIF operation. The user can specify the expected spectrum so that the output audio is reshaped to a flat, pink or voice-like curve. An adapt rate setting controls the update rate for the spectral average, which in turn determines how quickly the filter responds to changes in the input audio. Upper and lower limit controls allow the user to specify the range over which equalization is applied, and a mode setting controls whether frequencies outside the equalization range are attenuated or left unaffected. The amount of spectral correction is adjustable using the Filter Amount control. The user can enable the auto-gain functionality to ensure that the output audio level is maintained at approximately the same as the input audio level. If the user disables the auto-gain, an output gain slider is available to manually boost the level of the output signal.

As an aid to visualizing the filter operation, the user can view the input and output audio traces as well as the filter coefficient trace.

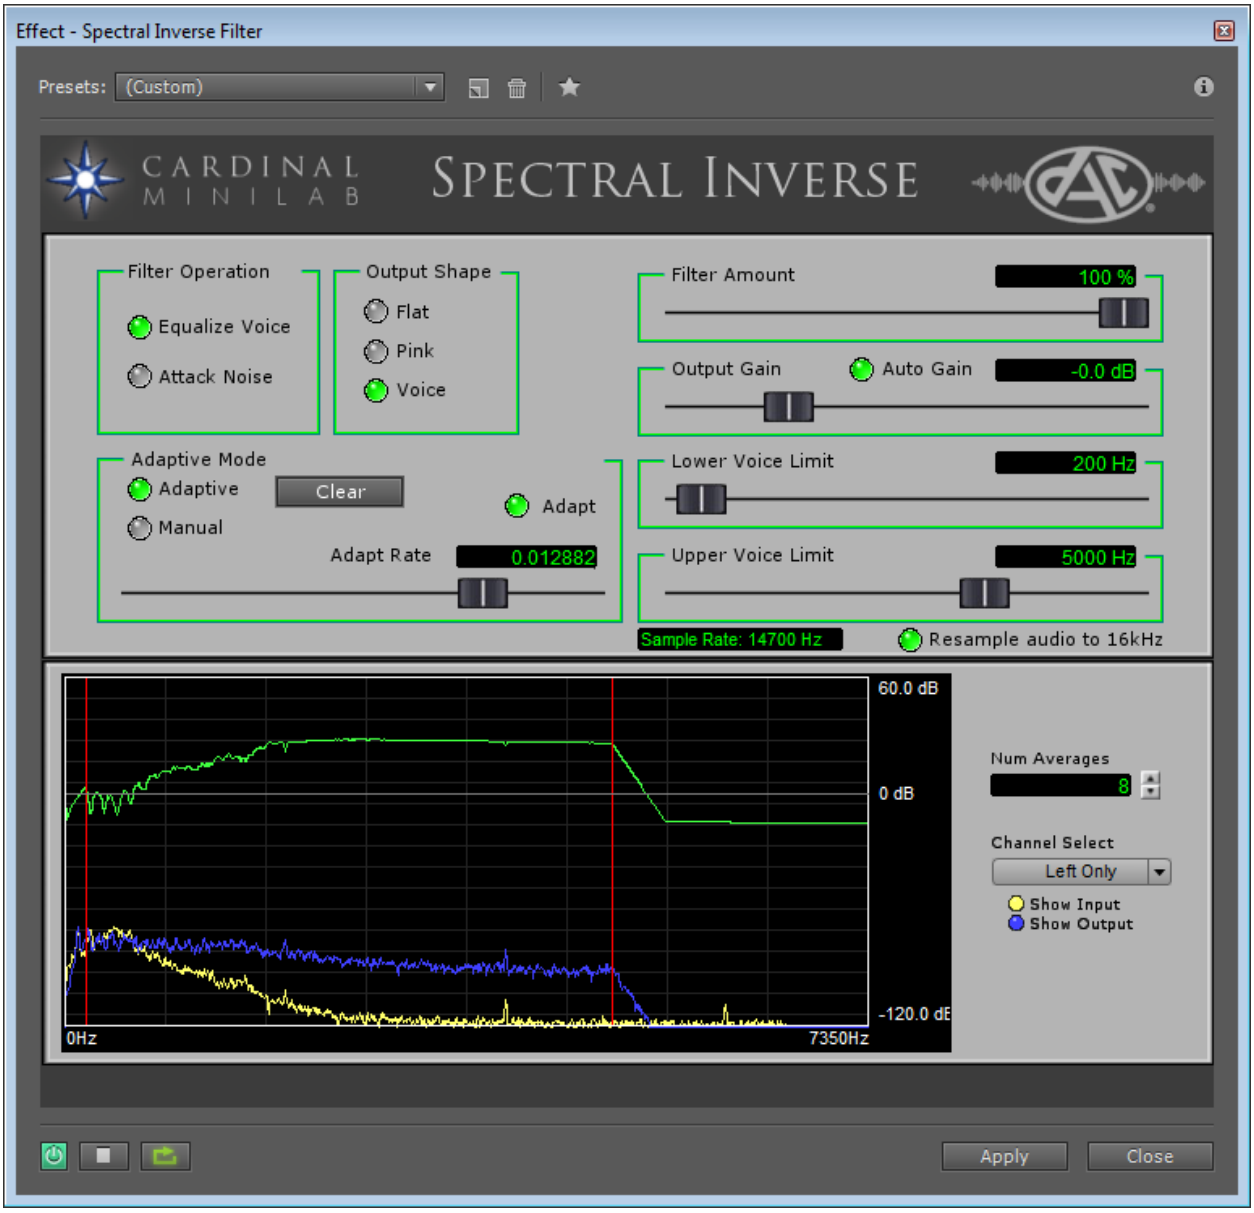

Figure 20: Spectral Inverse Filter Main Window (in Adaptive Mode)

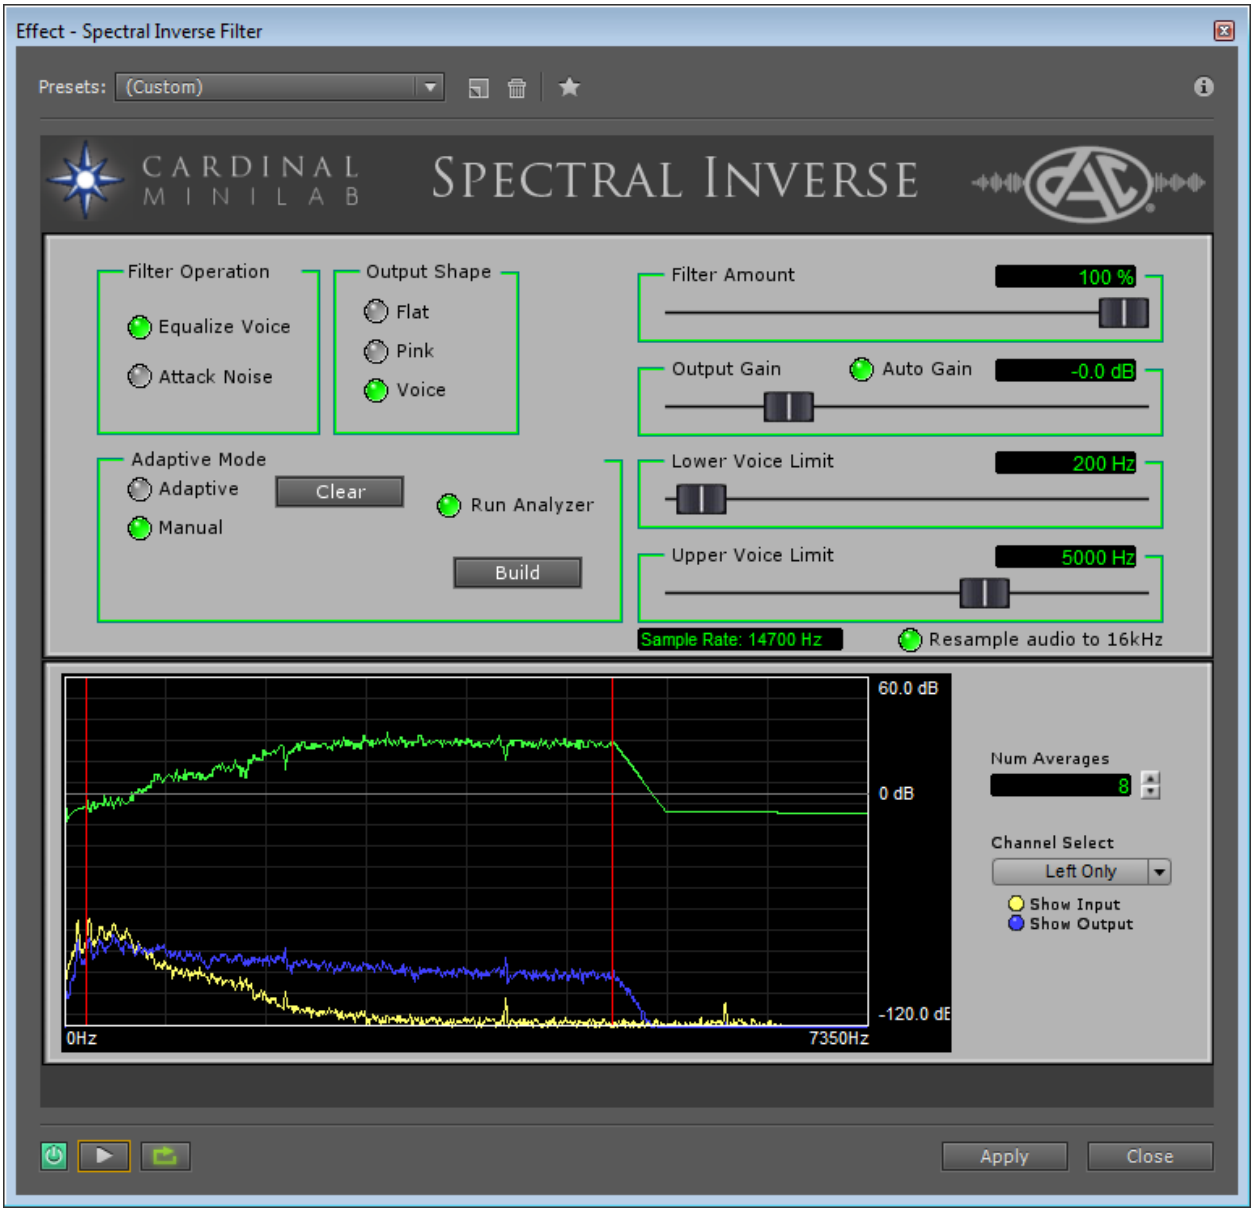

Figure 21: Spectral Inverse Filter Main Window (in Manual Mode)

6.2 MAIN PLUG-IN WINDOW

Figure 20 and Figure 21 show the plug-in as it would be displayed if used in Adobe Audition in **Adaptive** mode and **Manual** mode repsectively. The buttons above and below the plug-in window will appear differently depending upon your audio editing environment. Please consult the documentation on your specific audio editing environment for details of how to use these features.

This window provides access to all the functionality of the plug-in. A spectrum

analyzer is provided as an aid in determining the characteristics of the target audio. Details of the plug-in are described in the following sections.

6.3 FILTER CONTROLS

6.3.1 Filter Operation

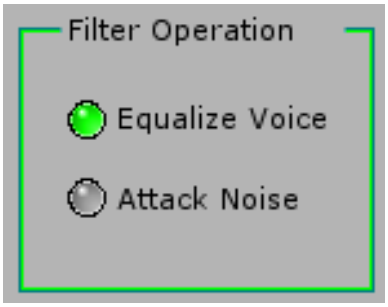

Figure 22: Filter Operation

In this block, the user can select the operational mode of the filter. If the filter is being used to correct spectral coloration, the **Equalize Voice** mode should be selected. If the filter is being used to remove bandlimited noise, the **Attack Noise** mode should be selected.

Note: The **Filter Operation** mode selection only affects the behavior of the filter *outside* the range selected by the upper and lower limits. In **Equalize Voice** mode, the frequency ranges outside the limits are attenuated. In **Attack Noise** mode, the frequency ranges outside the limits are left unaffected (subject to a transition region near the limits). If the auto gain is disabled and the manual gain is set to 0 dB, frequencies outside the limits and transition regions will be unaffected. However, if gain is applied, the gain will be reflected over the entire frequency range. See the section on **Upper** and **Lower Voice Limits** for more information on selecting the range.

Note: Changing the filter operation mode does not require an adaptation period to arrive at a "good" solution. Because a full average spectrum is maintained regardless of the mode setting, the new mode takes effect instantaneously in both the output audio and the display traces. However, since the auto gain adapts based on the actual applied filter with operational mode taken into account, there may be some adaptation time required to reach a stable auto gain value after the mode is changed.

6.3.2 Output Shape

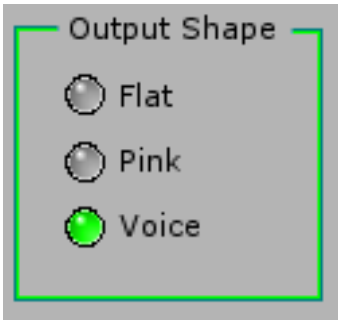

Figure 23: Output Shape

In this block, the user can select the target spectral shape that the filter attempts to achieve. The SIF has an inherent spectral flattening effect on the audio. The selected spectral shape is applied to further reshape the audio spectrum. The following output shapes are available:

- **Flat** – no additional shaping after ASIF flattening
- **Pink** -- 3 dB per octave rolloff above 100 Hz is applied in addition to ASIF flattening
- **Voice** – 6 dB/octave rolloff above and below 500 Hz in addition to ASIF flattening

Note: Changing the output shape does not require an adaptation period to arrive at a "good" solution. Because a full average spectrum is maintained regardless of the output shape setting, the new output shape takes effect instantaneously in both the output audio and the display traces. However, since the auto gain adapts based on the actual applied filter with the shaping curve taken into account, there may be some adaptation time required to reach a stable auto gain value after the shaping curve is changed.

6.3.3 Filter Amount

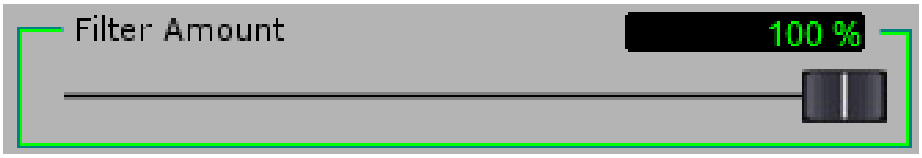

Figure 24: Filter Amount

This setting controls the degree to which the SIF can affect the signal, with 0% corresponding to no filtering and 100% corresponding to full filtering. In general, it is best to use the minimum Filter Amount setting that produces the desired result. When **Equalize Voice** mode is used, a lower **Filter Amount** can reduce artifacts that result from a fast adapt rate, so the **Filter Amount** can be used to help strike a balance between responsiveness and stability. When **Attack Noise** mode is used to reduce bandlimited noise, a lower **Filter**

Amount setting will often be a better choice to prevent the elevation of background noises.

Note: Changing the **Filter Amount** setting does not require an adaptation period to arrive at a "good" solution. Because a full average spectrum is maintained regardless of the setting, the new filter amount setting takes effect instantaneously in both the output audio and the display traces. However, since the auto gain adapts based on the actual applied filter with filter amount taken into account, there may be some adaptation time required to reach a stable auto gain value after the filter amount is adjusted.

6.3.4 Output Gain and Auto Gain

Figure 25: Output Gain and Auto Gain

These controls provide two options for adjusting the level of the SIF output. When **Auto Gain** is enabled, the SIF automatically monitors the input and output levels and applies a gain value that matches the output level to the input level. When **Auto Gain** is disabled, the user can use the **Output Gain** setting to specify the amount of boost applied to the SIF output. If **Auto Gain** is enabled, then the **Output Gain** slider is ignored.

The **Auto Gain** is an adaptive value whose rate of change depends on the same **Adapt Rate** slider setting that controls filter coefficient averaging. This means that when the filter response changes rapidly and dramatically, the auto gain will take some time to "catch up" to these changes. In particular, the output audio may clip when user settings are changed in a ways that have a boosting effect, such as switching from a pink to a flat shaping curve, adjusting the filter amount, or increasing the size of the ASIF region in Equalize Voice mode so that some frequencies that had been heavily attenuated are now present. While these settings changes will take effect immediately, the Auto Gain may take some time to adapt to the change. For this reason, when the user expects to be making many changes in the settings, it is often better to disable Auto Gain and instead choose a manual gain setting that avoids clipping.

6.3.5 Lower and Upper Voice Limits

Figure 26: Lower and Upper Voice Limit

The **Lower** and **Upper Voice Limits** allow the user to specify the frequency range, or "SIF region," over which the SIF is applied. Two red markers on the **Spectrum Analyzer** indicate where the lower and upper voice limits are located. Viewing audio on the display trace while manipulating the markers is an easy way to identify where your SIF region limits should fall.

In **Equalize Voice** mode, the SIF region is typically chosen to be the range over which speech frequencies are found. Setting a **Lower Limit** above 300 Hz or an **Upper Limit** below 3000 Hz is not recommended in equalize voice mode, as intelligibility may suffer. When in **Equalize Voice** mode, all frequencies outside the SIF region are assumed to be non-speech and are therefore attenuated.

In **Attack Noise** mode, the SIF region is typically chosen to "bracket" the bandlimited noise as closely as possible. Frequencies outside the ASIF region will be "passed through," i.e. there will be little or no effect outside the SIF region except for a narrow transition band between the SIF region and the passbands.

Note: Changing the Voice Limits does not require an adaptation period to arrive at a "good" solution. Because a full average spectrum is maintained regardless of the Voice Limit settings, the new Voice Limits will take effect instantaneously in both the output audio and the display traces. However, since the auto gain adapts based on the actual applied filter with voice limits taken into account, there may be some adaptation time required to reach a stable auto gain value after the limits are changed.

6.3.6 Adaptive Mode

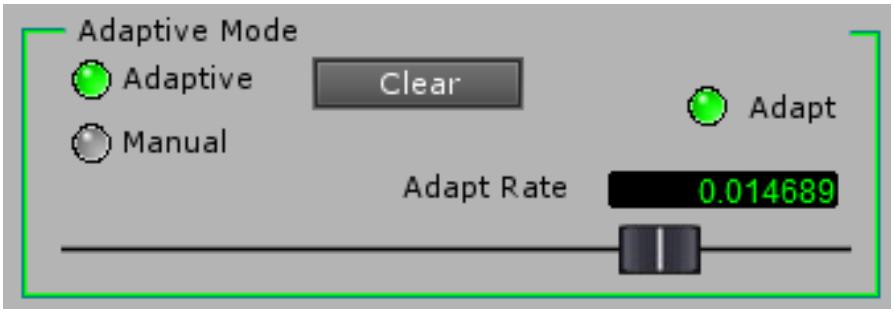

Figure 27: Adaptive Mode

In **Adaptive** mode, the above figure is displayed and gives the user access to the **Adapt, Clear** and **Adapt Rate** controls. This puts the plug-in in the ASIF mode.

When the **Adapt** button is lit green, the ASIF is adapting in response to incoming audio. When the button is grayed, the ASIF response is frozen.

The **Clear** button allows the user to re-initialize the ASIF response and restart adaptation.

Note: After a **Clear** operation or after re-enabling adaptation, there will be an adaptation period while the filter adapts to the current input signal. The length of this adaptation period depends on the **Adapt Rate** control setting.

The **Adapt Rate** control allows the user to select the rate of adaptation for the spectral average on which the ASIF response is based. The spectral averager uses an exponential average of the form $H_{i+1} = (\alpha)(X_{i+1}) + (1- \alpha)(H_i)$. The value shown in the display box corresponds to the averaging constant α in the exponential average. The lower the adapt rate value, the slower the filter will respond to changes in the input audio.

Note: "Fast response" sounds like a good thing, so it can be tempting to set the adapt rate to a high value. However, the goal of the ASIF is not to remove transient noises, but rather to reshape the long-term spectral envelope of the signal. If the adapt rate is too fast, the filter will respond too quickly to transient audio characteristics, which will produce artifacts in the output audio and will prevent the filter from settling on a good average solution. *For this reason, most applications will work best with adapt rates at the low end of the available range.* If you hear tonal artifacts that come and go in the output audio, or if the filter trace display coefficients seem to be changing rapidly, you probably need to reduce the adapt rate.

6.3.7 Manual Mode

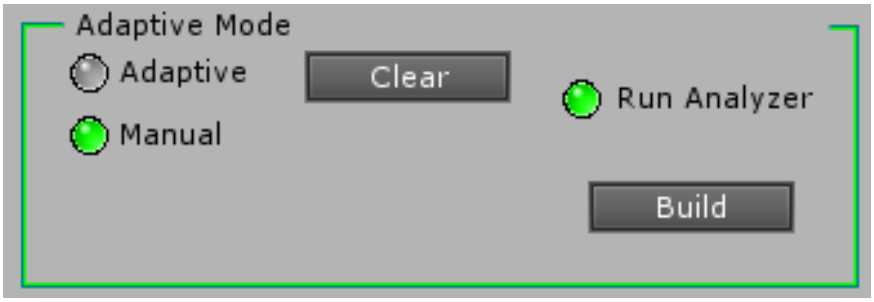

Figure 28: Manual Mode

In **Manual** mode, the above figure is displayed and gives the user access to the **Run Analyzer, Clear** and **Build** controls. This puts the plug-in in the SIF mode.

The **Clear** button is used to zero the averager memory and causes the averaged spectrum to be recalculated anew.

The **Run Analyzer** button allows the user to start or stop the update of the averaged spectrum.

The **Build** button builds the spectral inverse filter based on the original input audio spectrum and the SIF control settings. Once the filter build is complete the calculated spectral inverse filter curve will be displayed as a green trace in the Filter Display area.

Hint: Before clicking the **Build** button, it is recommended that the spectrum analyzer be stopped (using the **Run Analyzer** button) to allow experimentation with the control settings for the same input spectrum.

6.3.8 Display

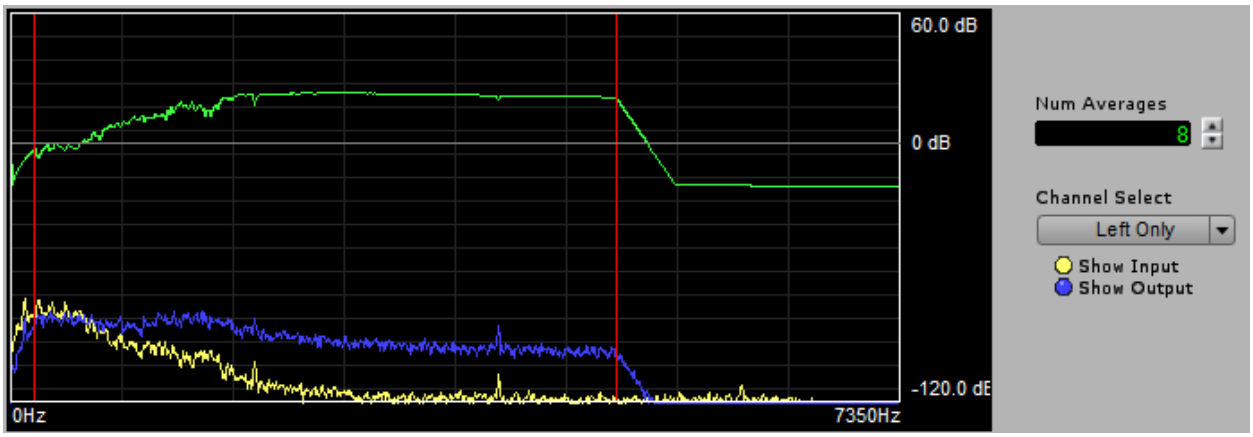

Figure 29: SIF Display

The display trace is used to view the filter input and output audio and the SIF filter response. The input audio is always shown in yellow, the output trace in blue and the filter trace in green. The **Lower** and **Upper Voice Limits** are shown in red.

The controls for this display are identical to the controls for the **Spectrum Analyzer**, except for the analyzer is always in the **Exponential** averaging mode (hence, no choice is given to change the averaging mode).

7 COMB FILTER

7.1 OVERVIEW OF THE COMB FILTER

The Comb filter is used to remove, or "notch out", harmonically related noises (noises which have exactly equally-spaced frequency components), such as power-line hum, constant-speed motor/generator noises, etc., from the input audio. The filter response consists of a series of equally-spaced notches which resemble a hair comb, hence the name "Comb filter".

Adjust the **Comb Frequency** to the desired spacing between notches (also known as "fundamental frequency"). Set the **Notch Limit** to the frequency beyond which you do not want any more notches. Set the **Notch Depth** to the amount in dB by which noise frequency components are to be reduced.

Normally, the Notch **Harmonics** option will be set to **All**, causing frequencies at all multiples of the Comb Frequency (within the Notch Limit) to be reduced. However, certain types of noises have only the odd or even harmonic components present. In these situations, set the Notch Harmonics option to either **Odd** or **Even**.

The **Auto Tracking** feature allows the filter to automatically hone in on the exact fundamental frequency and continually track it even if it shifts slightly.

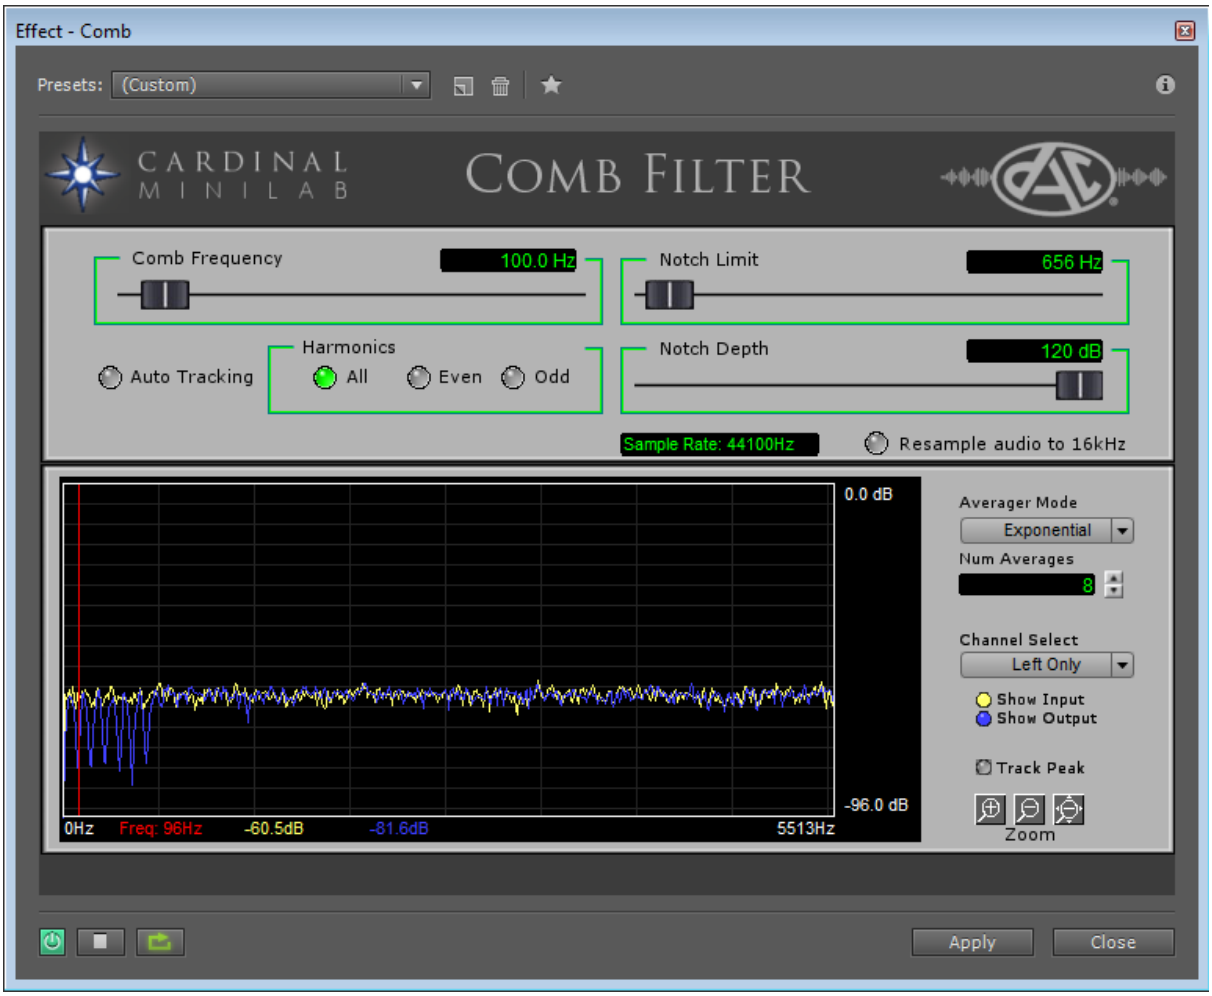

Figure 30: Comb Filter Main Window – Manual mode

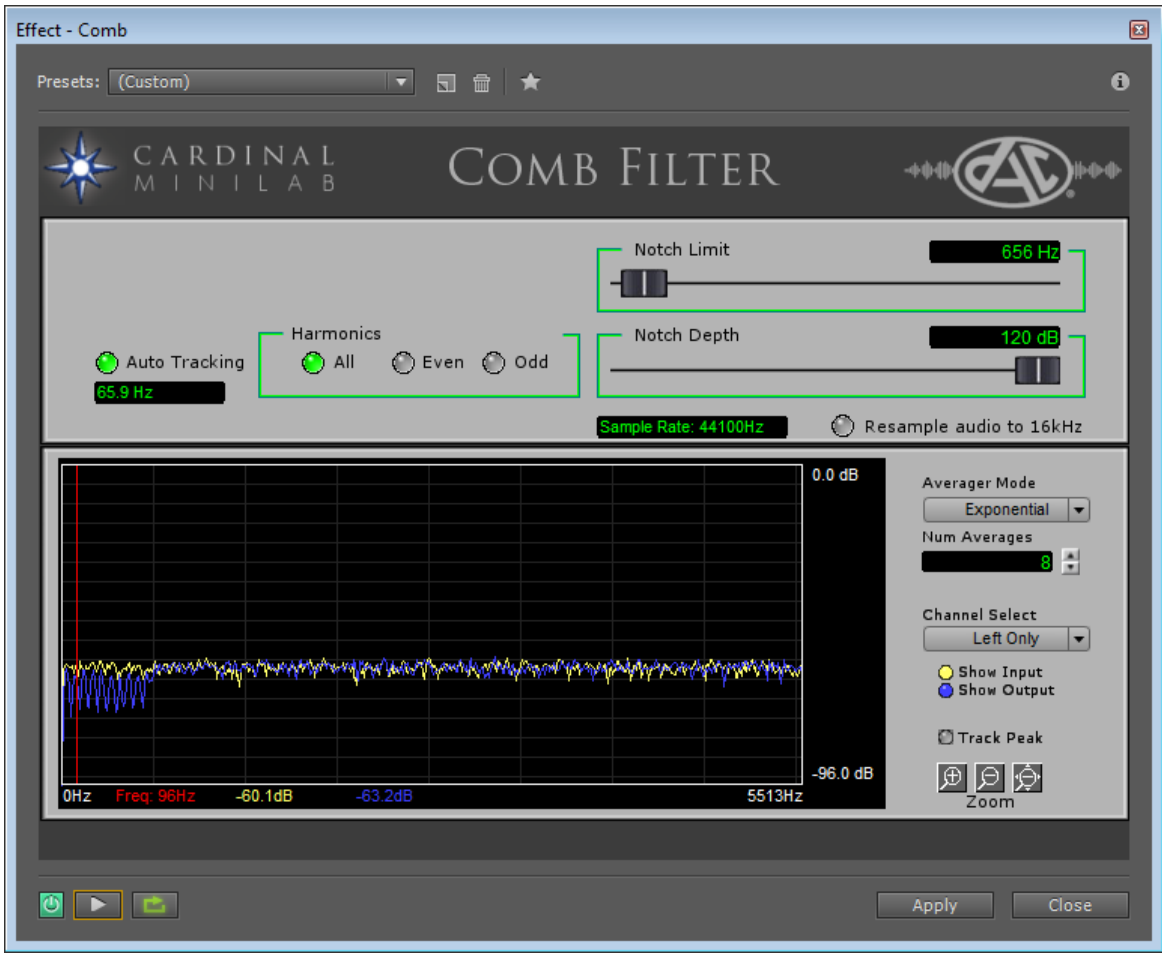

Figure 31: Comb Filter Main Window - Auto Tracking

7.1 MAIN PLUG-IN WINDOW

Figure 30 and Figure 31 above show the plug-in as it would be displayed if used in Adobe Audition in **Manual** mode and **Auto Tracking** mode respectively. The buttons above and below the plug-in window will appear differently depending upon your audio editing environment. Please consult the documentation on your specific audio editing environment for details of how to use these features.

This window provides access to all the functionality of the plug-in. A spectrum analyzer is provided as an aid in determining the characteristics of the target audio. Details of the plug-in are described in the following sections.

7.2 FILTER CONTROLS

7.2.1 Comb Frequency

Figure 32: Comb Frequency

Specifies fundamental frequency in Hertz of comb filter. Notches are generated at multiples, or harmonics, of this frequency. When **Auto Tracking** is turned on, this control does not appear.

7.2.2 Notch Limit

Figure 33: Notch Limit

Specifies frequency in Hertz above which no notches are generated. Minimum Notch Limit is half the **Comb Frequency**, while maximum Notch Limit depends upon the sample rate.

7.2.3 Notch Depth

Figure 34: Notch Depth

Specifies the depth of notches that are generated. Notch Depth is adjustable from 0 dB to 120 dB in 1 dB steps.

7.2.4 Notch Harmonics

Figure 35: Notch Harmonics

Specifies whether notches will be generated at All, Odd, or Even multiples, or harmonics, of the Comb Frequency. If, for example, the Comb Frequency is set to 60.0 Hz, then selecting **All** will generate notches at 60 Hz, 120 Hz, 180 Hz, 240 Hz, 300 Hz, etc. Selecting **Odd** will generate notches at 60 Hz, 180 Hz, 300 Hz, etc. Selecting **Even** will generate notches at 120 Hz, 240 Hz, 360 Hz etc.

7.2.5 Auto Tracking

Figure 36: Auto Tracking

Enables **Auto Tracking** for the Comb filter. When enabled, the **Comb Frequency** control is no longer available and the filter will determine the fundamental frequency to use.

When disabled, the last calculated frequency will then be loaded into the **Comb Frequency** control.

*NOTE: It is advised that you adjust the **Comb Frequency** manually to get it as close as possible to the desired fundamental frequency before enabling **Auto Tracking**.*

8 X-PASS FILTER

8.1 OVERVIEW OF THE X-PASS FILTER

The X-Pass Filter provides six different bandlimiting filters in a single plug-in. The following filters are included:

Filter	Application
Lowpass	Decrease the energy level of high frequency noises
Highpass	Decrease the energy level of low frequency noises
Bandpass	Decrease the energy level of noises above and below a passband region
Bandstop	Decrease the energy level of noises within a stopband region
Notch	Remove a narrow-band noise
Slot	Isolate a narrow-band signal

Only one filter can be operational at a time. While they all share similar controls, they perform different functions (as detailed below). Whichever filter you select, its transfer function (or envelope) is displayed in the Spectrum Analyzer as a light gray trace.

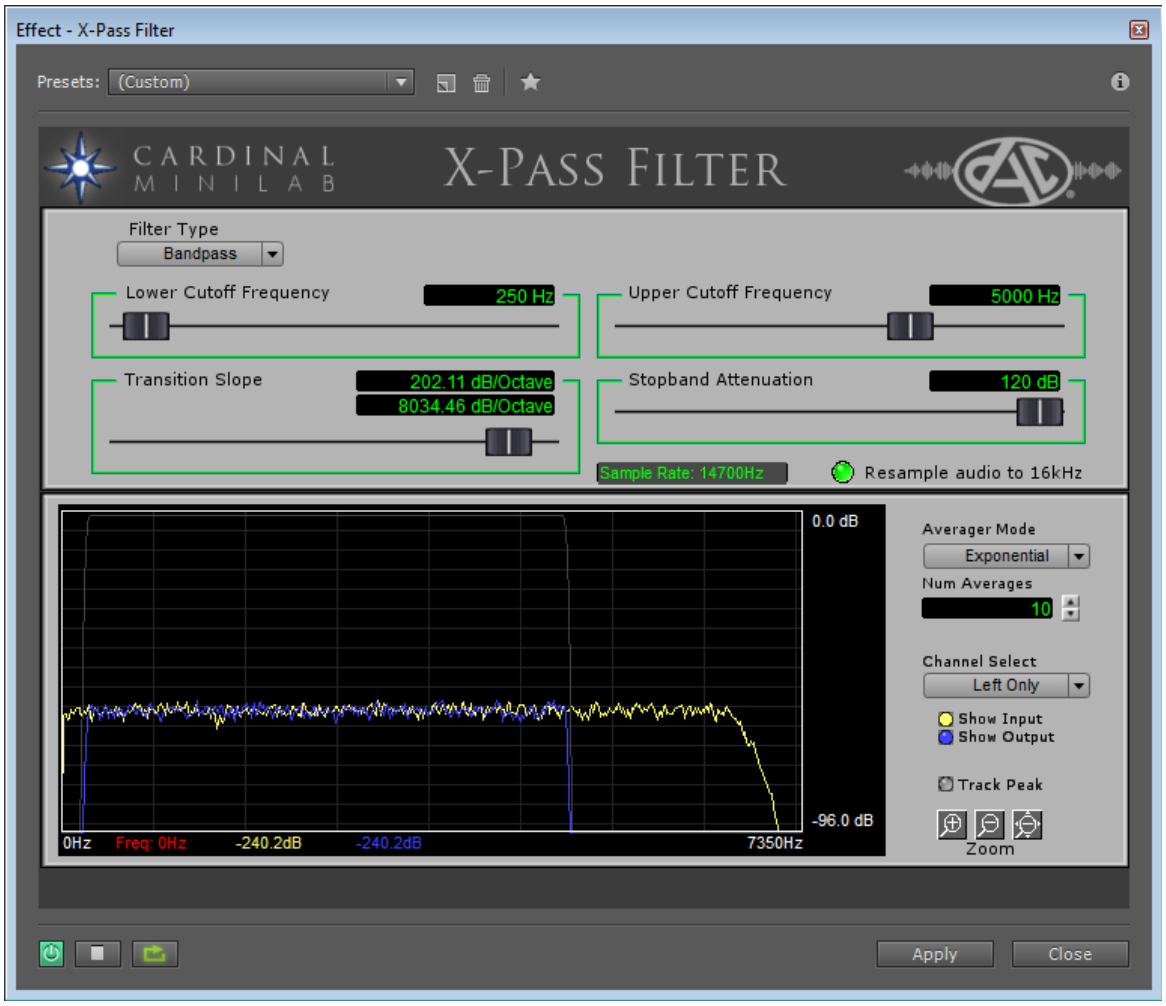

Figure 37: X-Pass Main Window

8.2 MAIN PLUG-IN WINDOW

Figure 37 above shows the plug-in as it would be displayed if used in Adobe Audition. The buttons above and below the plug-in window will appear differently depending upon your audio editing environment. Please consult the documentation on your specific audio editing environment for details of how to use these features.

The **Filter Type** control selects which of the six filters is currently engaged. Only one filter is operational at a time. If you desire to apply more than one of the filters on a given piece of audio, you must either use multiple instances of the plug-in in an effects chainer (not available in all audio editor environments) or render the audio multiple times (using a different filter each time). When a filter is selected, that filter's controls are displayed in the main window.

This window provides access to all the functionality of the plug-in. A spectrum analyzer is provided as an aid in determining the characteristics of the target audio. Details of the plug-in are described in the following sections.

8.3 LOWPASS FILTER

The Lowpass filter is used to decrease the energy level (lower the volume) of all signal frequencies above a specified Cutoff Frequency, thus reducing high-frequency noises, such as tape hiss, from the input audio. The Lowpass filter is sometimes called a "hiss filter."

The Cutoff Frequency is usually set above the voice frequency range so that the voice signal will not be disturbed. While listening to the filter output audio, the Cutoff Frequency can be incrementally lowered from its maximum frequency until the quality of the voice just begins to be affected, achieving maximum elimination of high-frequency noise.

The amount of volume reduction above the Cutoff Frequency can further be controlled by adjusting the Stopband Attenuation setting (maximum volume reduction is 120dB). The slope at which the volume is reduced from normal (at the Cutoff Frequency) to the minimum volume (specified by Stopband Attenuation) can also be controlled by adjusting the Transition Slope setting.

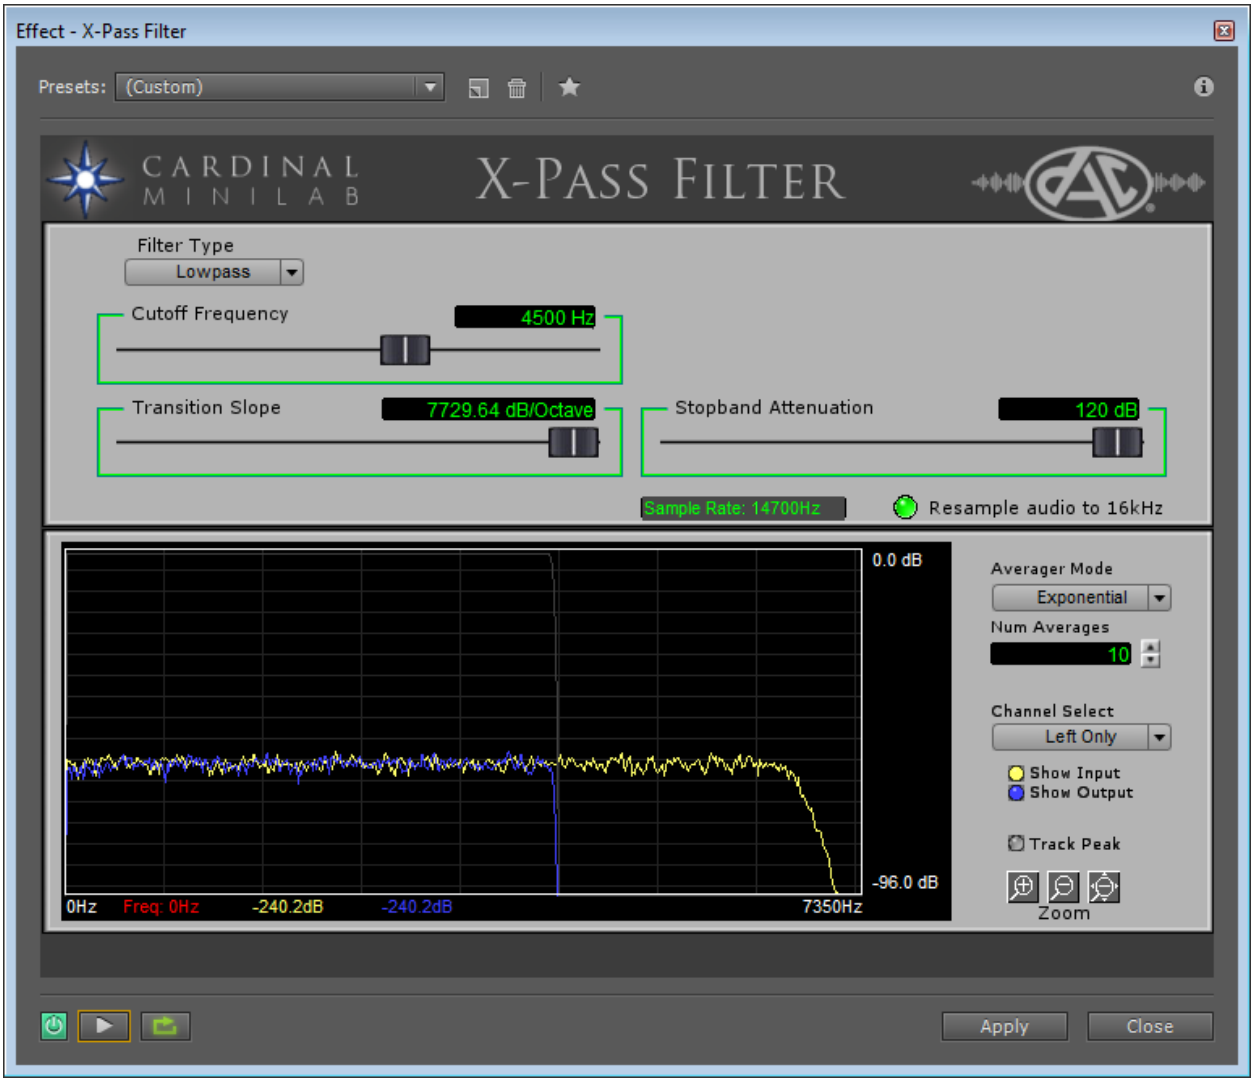

Figure 38: Lowpass Filter Controls

8.3.1 Filter Controls

8.3.1.1 Cutoff Frequency

Figure 39: Cutoff Frequency

Specifies frequency in Hertz above which all signals are attenuated. Frequencies below this cutoff are unaffected. Maximum Cutoff Frequency depends upon the Sample Rate. Cutoff Frequency can be adjusted in 1 Hz steps.

8.3.1.2 Transition Slope

Figure 40: Transition Slope

Specifies slope at which frequencies above the Cutoff Frequency are rolled off in dB per octave. Sharpest roll off occurs when Transition Slope is set to maximum, while gentlest roll off occurs when Transition Slope is set to minimum. Sharp rolloffs may cause the voice to sound hollow but will allow more precise removal of high frequency noises. Note that the indicated value changes depending upon the Cutoff Frequency and Sample Rate.

8.3.1.3 Stopband Attenuation

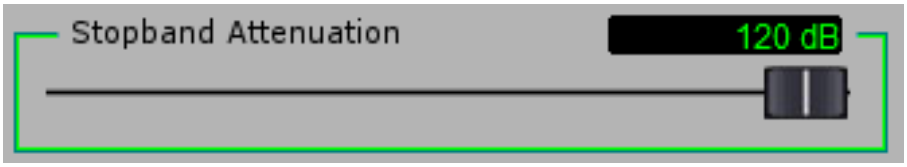

Figure 41: Stopband Attenuation

Specifies amount in dB by which frequencies above the Cutoff Frequency are ultimately attenuated. Stopband attenuation is adjustable from 0dB to 120dB in 1 dB steps.

8.4 HIGHPASS FILTER

The Highpass filter is used to decrease the energy level (lower the volume) of all signal frequencies below a specified Cutoff Frequency, thus reducing low-frequency noises, such as tape or acoustic room rumble, from the input audio (The Highpass filter is sometimes called a "rumble filter").

The Cutoff Frequency is usually set below the voice frequency range (somewhere below 300 Hz) so that the voice signal will not be disturbed. While listening to the filter output audio, the Cutoff Frequency, initially set to 0 Hz, can be incrementally increased until the quality of the voice just begins to be affected, achieving maximum elimination of low-frequency noise.

The amount of volume reduction below the Cutoff Frequency can further be controlled by adjusting the Stopband Attenuation setting (maximum volume reduction is 120dB). The slope at which the volume is reduced from normal (at the Cutoff Frequency) to the minimum volume (specified by Stopband

Attenuation) can also be controlled by adjusting the Transition Slope setting.

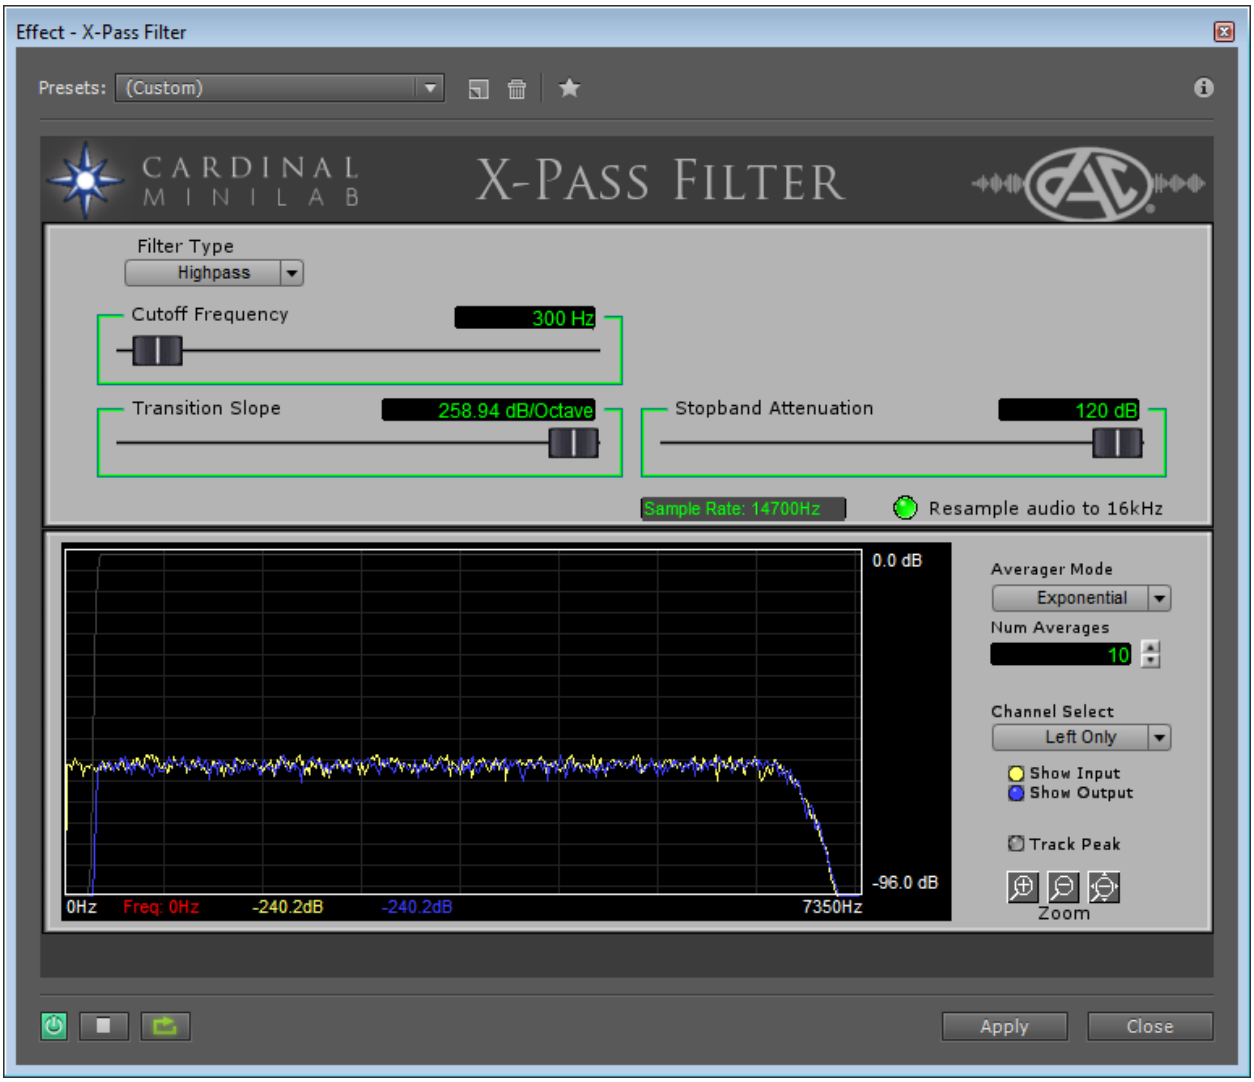

Figure 42: HIghpass Filter Controls

8.4.1 Filter Controls

8.4.1.1 Cutoff Frequency

Figure 43: Cutoff Frequency

Specifies frequency in Hertz below which all signals are attenuated. Frequencies

above this cutoff are unaffected. Minimum Cutoff Frequency is 0 Hz (no frequencies attenuated), while the maximum Cutoff Frequency depends upon the Sample Rate. Cutoff Frequency can be adjusted in 1 Hz steps.

8.4.1.2 Transition Slope

Figure 44: Transition Slope

Specifies slope at which frequencies below the Cutoff Frequency are attenuated in dB per octave. Sharpest attenuation occurs when Transition Slope is set to maximum, while gentlest attenuation occurs when Transition Slope is set to minimum. Note that the indicated value changes depending upon the Cutoff Frequency and Sample Rates.

8.4.1.3 Stopband Attenuation

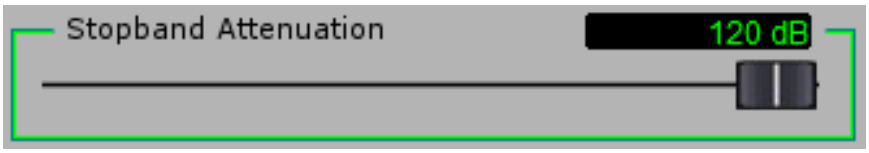

Figure 45: Stopband Attenuation

Specifies amount in dB by which frequencies below the Cutoff Frequency are ultimately attenuated.

8.5 BANDPASS FILTER

The Bandpass filter is used to decrease the energy level (lower the volume) of all signal frequencies below a specified Lower Cutoff Frequency and above a specified Upper Cutoff Frequency, thus combining the functions of a Lowpass and Highpass filter connected in series into a single filter. The signal region between the Lower Cutoff Frequency and the Upper Cutoff Frequency is called the passband region. The Bandpass filter is useful for simultaneously reducing both low-frequency rumble and high-frequency hiss.

The Lower Cutoff Frequency is usually set below the voice frequency range (somewhere below 300 Hz) so that the voice signal will not be disturbed. While listening to the filter output audio, the Lower Cutoff Frequency, initially set to 0 Hz, can be incrementally increased until the quality of the voice just begins to be affected, achieving maximum elimination of low-frequency noise.

The Upper Cutoff Frequency is usually set above the voice frequency range (somewhere above 3000 Hz) so that the voice signal will not be disturbed. While listening to the filter output audio, the Upper Cutoff Frequency, initially set to its maximum frequency, can be incrementally lowered until the quality of the voice just begins to be affected, achieving maximum elimination of high-frequency noise.

The amount of volume reduction outside the passband region can further be controlled by adjusting the Stopband Attenuation setting (maximum volume reduction is 120dB). The slope at which the volume is reduced from normal (at each Cutoff Frequency) to the minimum volume (specified by Stopband Attenuation) can also be controlled by adjusting the Transition Slope setting.

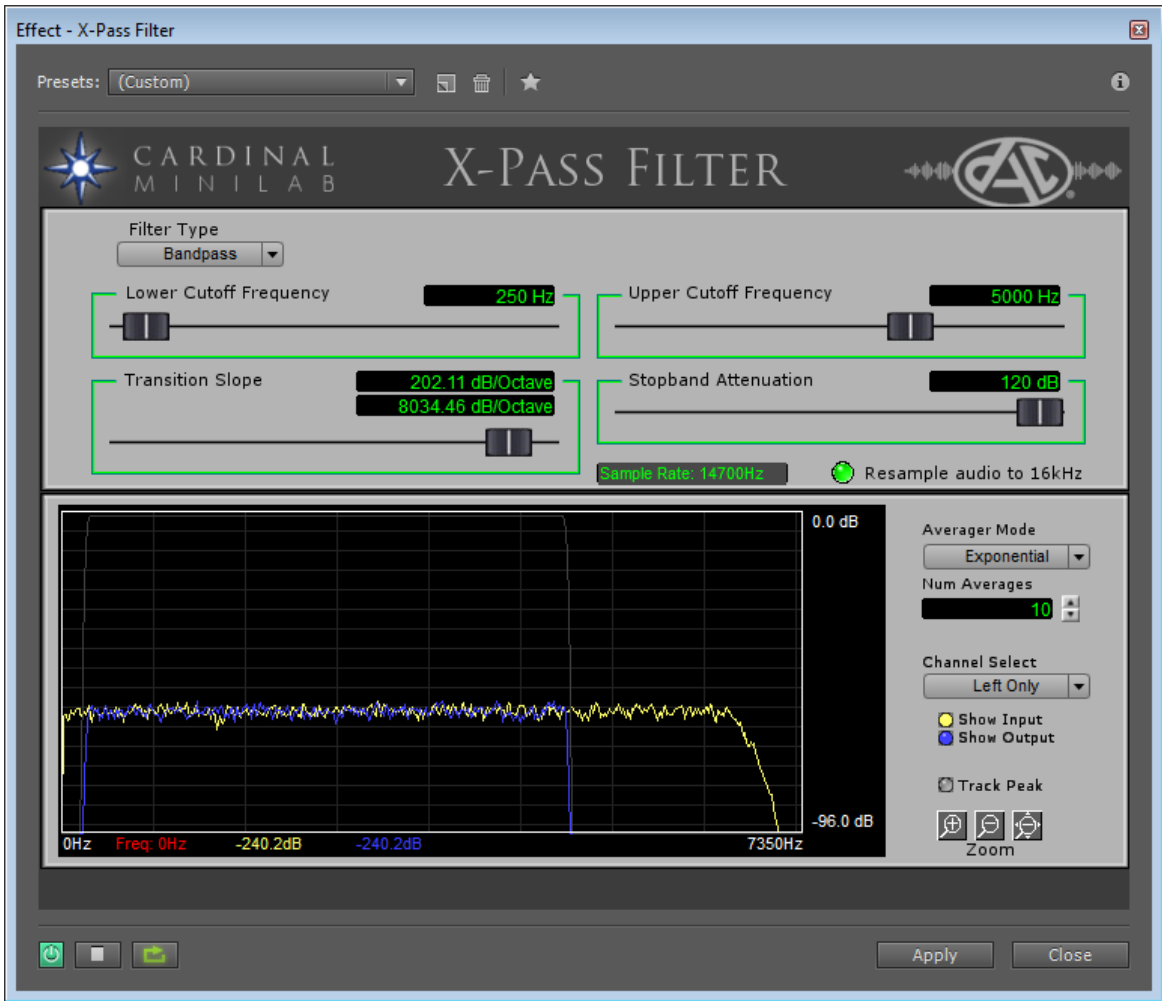

Figure 46: Bandpass Filter Controls

8.5.1 Filter Controls

8.5.1.1 Lower Cutoff Frequency

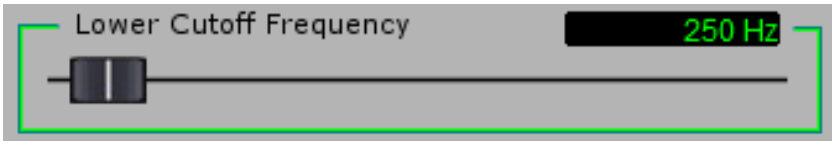

Figure 47: Lower Cutoff Frequency

Specifies frequency in Hertz below which all signals are attenuated. Frequencies between this cutoff and the Upper Cutoff Frequency are unaffected. Minimum Lower Cutoff Frequency is 0 Hz, while the maximum Lower Cutoff Frequency is 10 Hz below the Upper Cutoff Frequency. Lower Cutoff Frequency can be adjusted in 1 Hz steps.

NOTE: The Lower Cutoff Frequency can never be set higher than 10 Hz below the Upper Cutoff Frequency.

8.5.1.2 Upper Cutoff Frequency

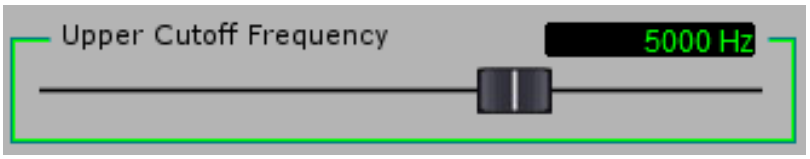

Figure 48: Upper Cutoff Frequency

Specifies frequency in Hertz above which all signals are attenuated. Frequencies between this cutoff and the Lower Cutoff Frequency are unaffected. Minimum Upper Cutoff Frequency is 10 Hz above the Lower Cutoff Frequency, while the maximum Upper Cutoff Frequency depends upon the Sample Rate. Upper Cutoff Frequency can be adjusted in 1 Hz steps.

NOTE: The Upper Cutoff Frequency can never be set lower than 10 Hz above the Lower Cutoff Frequency.

8.5.1.3 Transition Slope

Figure 49: Transition Slope

Specifies slope at which frequencies below the Lower Cutoff Frequency and above the Upper Cutoff Frequency are attenuated in dB per octave. Sharpest attenuation occurs when Transition Slope is set to maximum, while gentlest

attenuation occurs when Transition Slope is set to minimum. Note that the indicated value changes depending upon the Cutoff Frequency and Sample Rates. Also, note that the Lower and Upper Transition Slopes always have different values; this is because the frequency width of an octave is proportional to Cutoff Frequency.

8.5.1.4 Stopband Attenuation

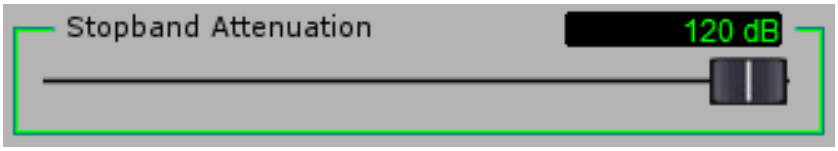

Figure 50: Stopband Attenuation

Specifies amount in dB by which frequencies below the Lower Cutoff Frequency and above the Upper Cutoff Frequency are ultimately attenuated.

8.6 BANDSTOP FILTER

The Bandstop filter is used to decrease the energy level (lower the volume) of all signal frequencies above a specified Lower Cutoff Frequency and below a specified Upper Cutoff Frequency. The signal region between the Lower Cutoff Frequency and the Upper Cutoff Frequency is called the stopband region. The Bandstop filter is useful for removing in-band noise from the input signal.

The Lower Cutoff Frequency is usually set below the frequency range of the noise, while the Upper Cutoff Frequency is set above the frequency range of the noise. While listening to the filter output audio, the Lower and Upper Cutoff Frequencies can be incrementally adjusted to achieve maximum elimination of noise while minimizing loss of voice.

The amount of volume reduction in the stopband region can further be controlled by adjusting the Stopband Attenuation setting (maximum volume reduction is 120dB). The slope at which the volume is reduced from normal (at each Cutoff Frequency) to the minimum volume (specified by Stopband Attenuation) can also be controlled by adjusting the Transition Slope setting.

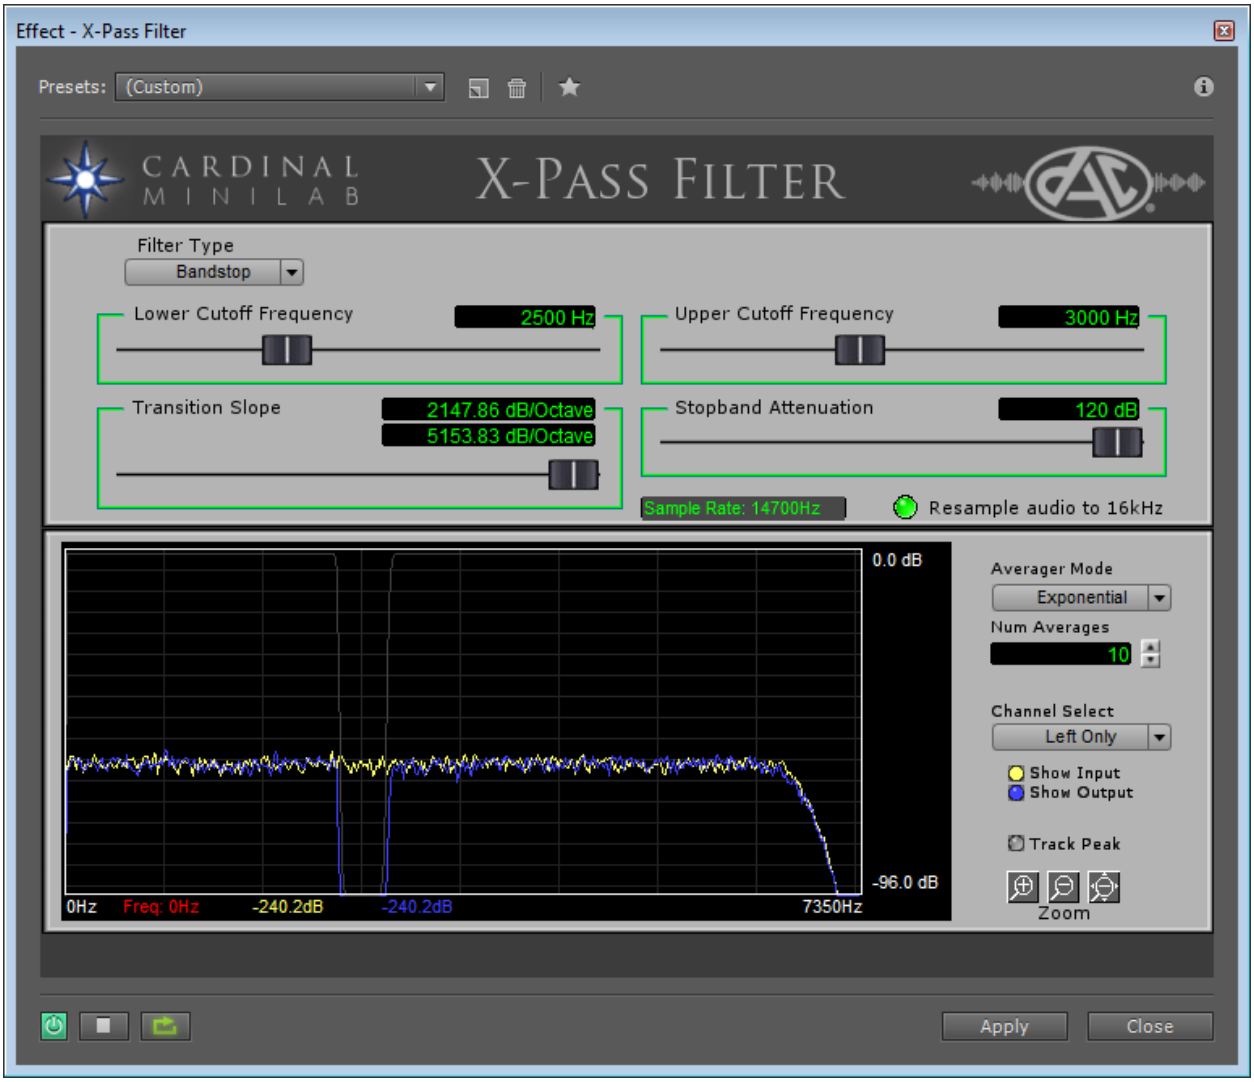

Figure 51: Bandstop Filter Controls

8.6.1 Filter Controls

8.6.1.1 Lower Cutoff Frequency

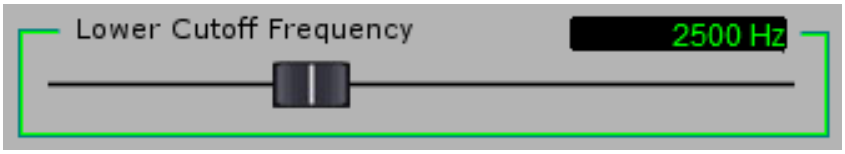

Figure 52: Lower Cutoff Frequency

Specifies frequency in Hertz below which no signals are attenuated. Frequencies between this cutoff and the Upper Cutoff Frequency are attenuated. Minimum Lower Cutoff Frequency is 0 Hz, while the maximum Lower Cutoff Frequency is

10 Hz below the Upper Cutoff Frequency. Lower Cutoff Frequency can be adjusted in 1 Hz steps.

NOTE: The Lower Cutoff Frequency can never be set higher than 10 Hz below the Upper Cutoff Frequency.

8.6.1.2 Upper Cutoff Frequency

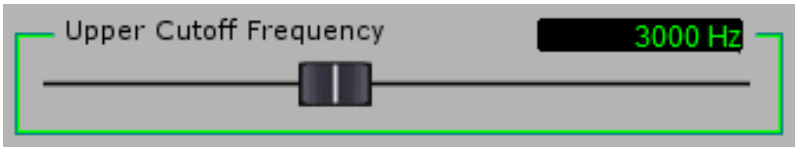

Figure 53: Upper Cutoff Frequency

Specifies frequency in Hertz above which no signals are attenuated. Frequencies between this cutoff and the Lower Cutoff Frequency are attenuated. Minimum Upper Cutoff Frequency is 10 Hz above the Lower Cutoff Frequency, while the maximum Upper Cutoff Frequency depends upon the Sample Rate. Upper Cutoff Frequency can be adjusted in 1 Hz steps.

NOTE: The Upper Cutoff Frequency can never be set lower than 10 Hz above the Lower Cutoff Frequency.

8.6.1.3 Transition Slope

Figure 54: Transition Slope

Specifies slope at which frequencies above the Lower Cutoff Frequency and below the Upper Cutoff Frequency are attenuated in dB per octave. Sharpest attenuation occurs when Transition Slope is set to maximum, while gentlest attenuation occurs when Transition Slope is set to minimum. Note that the indicated value changes depending upon the Cutoff Frequency and Sample Rates. Also, note that the Lower and Upper Transition Slopes always have different values; this is because the frequency width of an octave is proportional to Cutoff Frequency.

8.6.1.4 Stopband Attenuation

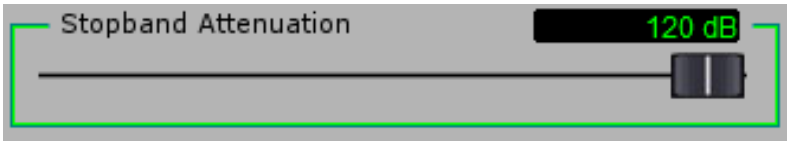

Figure 55: Stopband Attenuation

Specifies amount in dB by which frequencies above the Lower Cutoff Frequency and below the Upper Cutoff Frequency are attenuated.

8.7 NOTCH FILTER

The Notch filter is used to remove, or "notch out", a narrow-band noise, such as a tone or a whistle, from the input audio with minimal effect to the remaining audio. The Notch filter works best with stable noise sources which have constant frequency; if the frequency of the noise source varies, the 1-Channel Adaptive filter is recommended.

To properly utilize the Notch filter, you will first need to identify the frequency of the noise; this is best done using the Spectrum Analyzer window.

Initially set the Notch Depth to 120 dB and the Notch Width to the narrowest possible value. Next, set the Notch Frequency to the noise frequency. Fine adjustment of the Notch Frequency may be necessary to place the notch precisely on top of the noise signal and achieve maximum reduction of the noise. This is best done by adjusting the Notch Frequency up or down 1 Hz at a time while listening to the Notch filter output on the headphones.

Often, the noise frequency will not remain absolutely constant but will vary slightly due to modulation, recorder wow and flutter, and acoustic "beating." Therefore, you may need to increase the Notch Width from its minimum setting to keep the noise within the notch.

For maximum noise reduction, set the Notch Depth to 120dB. It is best to adjust the Notch Depth up from 120 dB until the tone is observed, then increase the depth 5 dB.

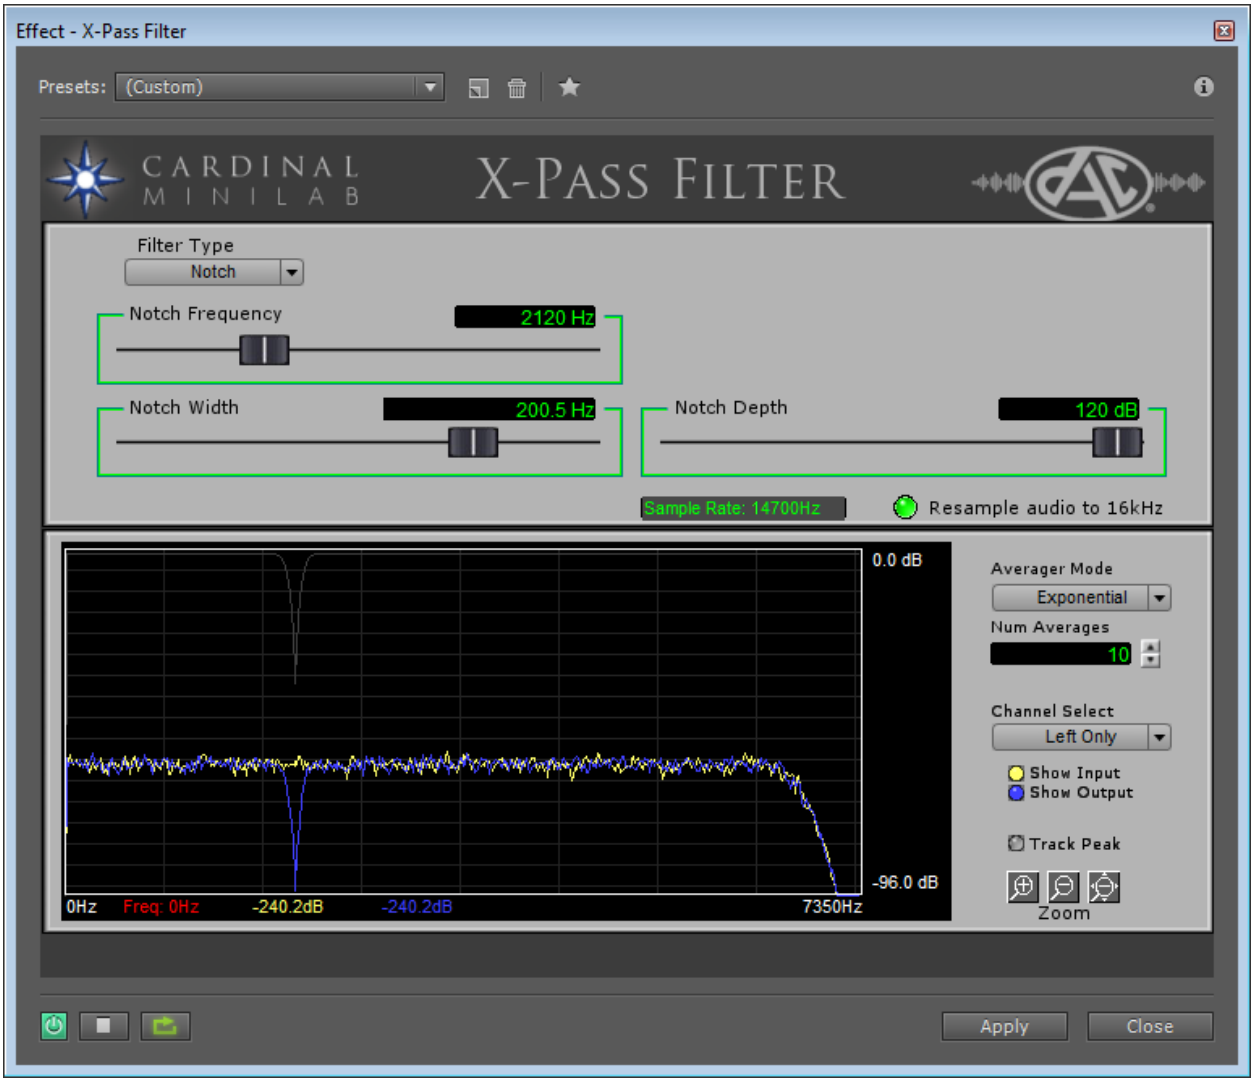

Figure 56: Notch Filter Controls

8.7.1 Filter Controls

8.7.1.1 Notch Frequency

Figure 57: Notch Frequency

Specifies frequency in Hertz which is to be removed from the input audio. Minimum Notch Frequency is 10 Hz, while maximum Notch Frequency depends upon the Sample Rate. Notch Frequency is adjustable in 1 Hz steps.

8.7.1.2 Notch Width

Figure 58: Notch Width

Width of the generated notch in Hertz.

NOTE: Notch Width varies with the Sample Rate.

8.7.1.3 Notch Depth

Figure 59: Notch Depth

Depth of the notch that is generated.

8.8 SLOT FILTER

The Slot filter is used to isolate, or "slot", a single-frequency signal, such as a tone or a whistle, in the input audio, attenuating all other audio. This is the exact opposite of the Notch filter function.

NOTE: The Slot filter has very little use in speech enhancement applications; the main value is in isolating other types of signals that are non-speech in nature.

To properly utilize the Slot filter, you will first need to identify the frequency of the signal to be isolated; this is best done using the Spectrum Analyzer window.

Once the frequency of the signal has been identified, initially set Stopband Attenuation to 120 dB and the Slot Width to the narrowest possible value. Next, set the Slot Frequency to the signal frequency. Fine adjustment of the Slot Frequency may be necessary to place the slot right on top of the signal. This is best done by adjusting the Slot Frequency up or down 1 Hz at a time while listening to the Slot filter output on the headphones.

Usually, the signal frequency will not remain constant but will vary slightly due to modulation, recorder wow and flutter, and acoustic "beating". Therefore, you

may need to increase the Slot Width from its minimum setting to avoid having the signal move in and out of the slot.

To optimize background noise reduction for your application, set the Stopband Attenuation to 120dB. If, however, you wish to leave a small amount of the background noise mixed in with the isolated signal, adjust the Stopband Attenuation to the desired value.

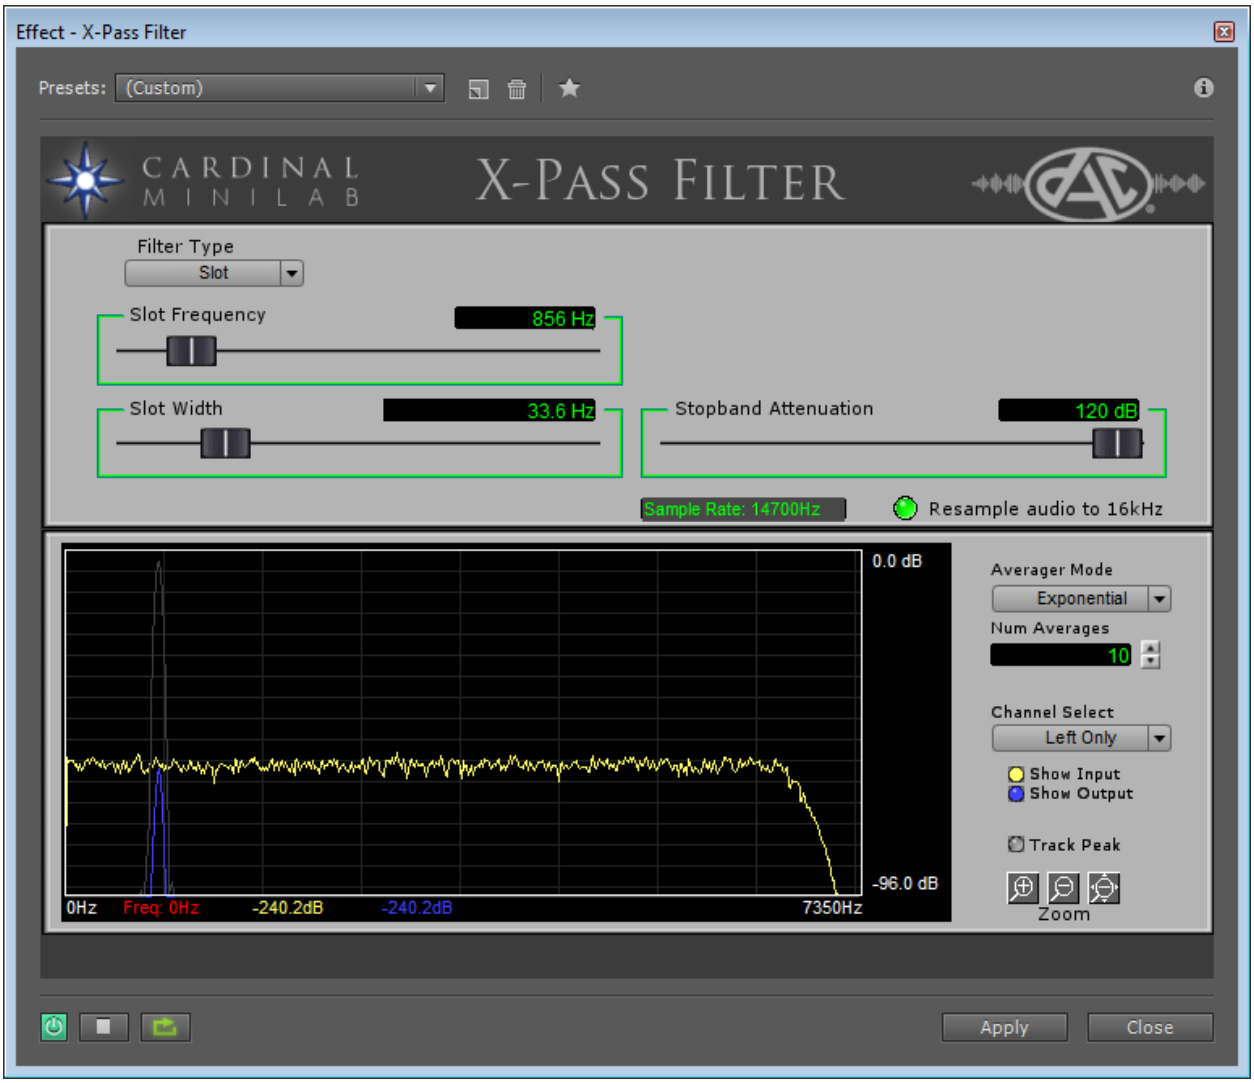

Figure 60: Slot Filter Controls

8.8.1 Filter Controls

8.8.1.1 Slot Frequency

Figure 61: Slot Frequency

Specifies frequency in Hertz which is to be enhanced in the input audio. Minimum Slot Frequency is 10 Hz, while maximum Slot Frequency depends upon the Sample Rate. Slot Frequency is adjustable in 1 Hz steps.

8.8.1.2 Slot Width

Figure 62: Slot Width

Width of the generated slot in Hertz.

NOTE: Slot Width varies with the Sample Rate.

8.8.1.3 Stopband Attenuation

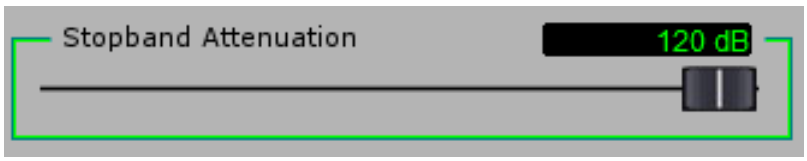

Figure 63: Stopband Attenuation

Specifies amount in dB by which frequencies other than the Slot Frequency are attenuated.

9 SPECTRAL SUBTRACTION FILTER

9.1 OVERVIEW OF THE SPECTRAL SUBTRACTION FILTER

The Spectral Subtraction filter is a frequency-domain filter that implements automatic noise reduction over the entire frequency spectrum. It operates by continually measuring the spectrum of the input signal and attempting to identify which portions of the signal are voice and which portions are non-voice (or noise). All portions determined to be noise are used to continually update a **noise estimate** calculation; this is used to calculate the equalization curve that needs to be applied to the input signal to reduce each band's energy by the amount of noise energy calculated to be in that band.

The net result is an output signal that has all non-voice signals reduced in level as much as possible, thereby "polishing" the enhanced voice signal as much as possible prior to final equalization and AGC.

The Spectral Subtraction filter has two modes: the **Noise Reducer** mode and the **Noise EQ** mode.

Operation in the Noise EQ mode is governed by 20 control sliders, each representing a frequency band. Adjusting the control sliders allows the user to precisely control the amount of noise reduction being applied within each of 20 distinct groups of frequency bands, offering much more precise control of the spectral subtraction than is available in the Noise Reducer mode, though it does take more time to setup.

In the Noise Reducer mode all 20 control sliders are linked together and are adjusted using the first slider. Adjusting the master slider allows the user to precisely control the amount of noise reduction being applied to all frequency bands; the greater the value, the more aggressive the operation of the Noise Reducer.

Because large amounts of noise reduction invariably create audible "birdy noise" artifacts in the output audio due to the nature of adaptive frequency-domain processing, the user should always try to minimize the amount of noise reduction being applied to achieve the best balance between maximal noise reduction and minimal audible artifacts.

Finally, for convenience an Output Gain control and Output level bargraph are provided to enable the user to adjust the processed output signal to maximum level for better listening and recording.

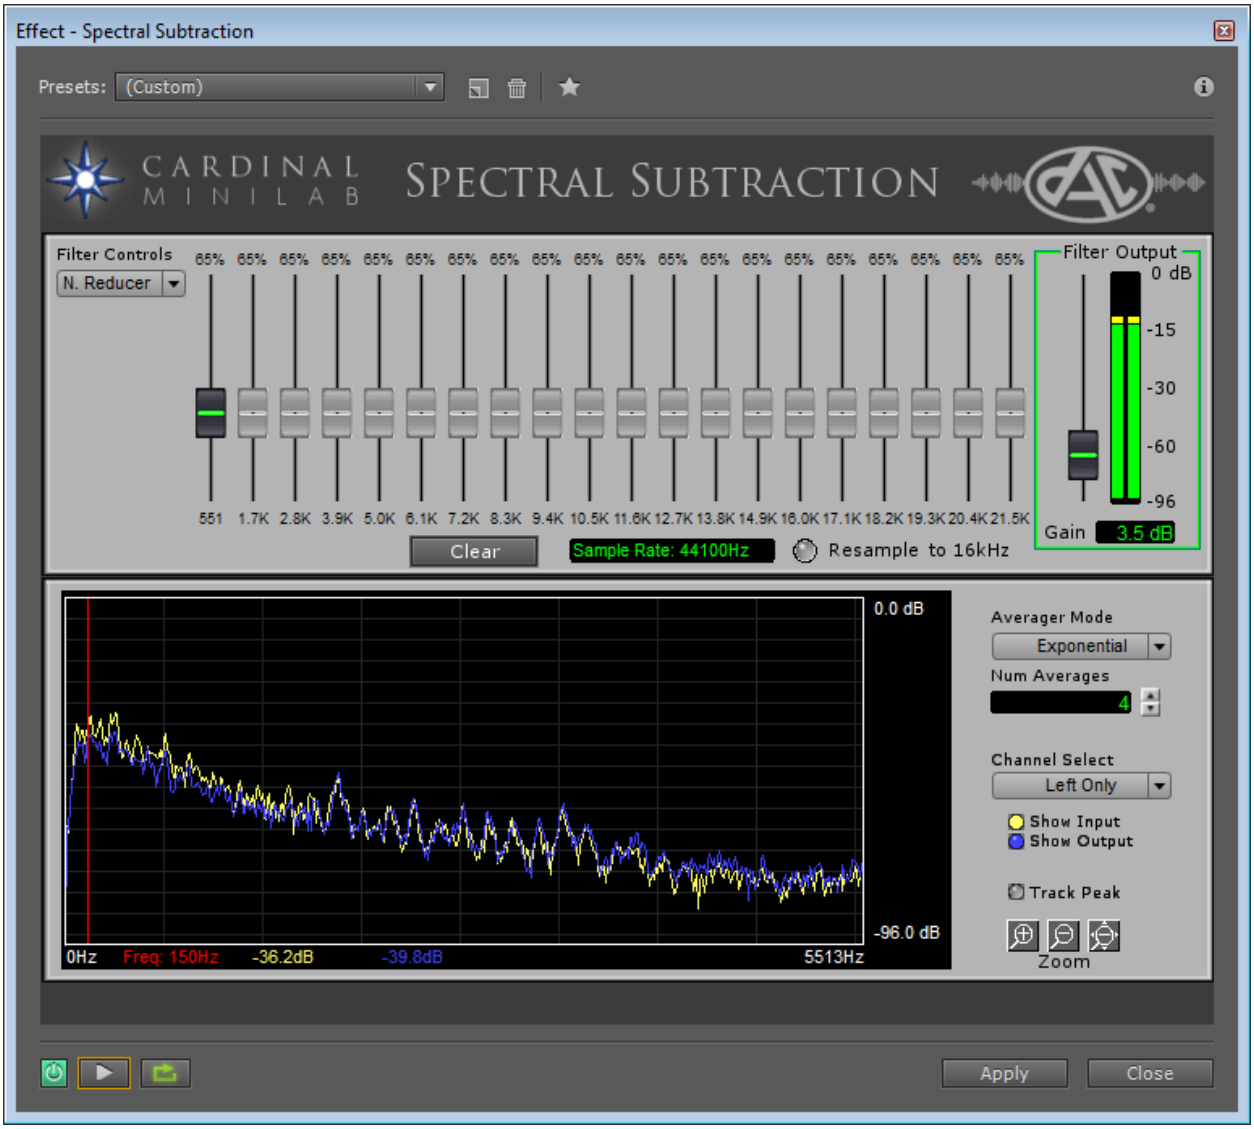

Figure 64: Spectral Subtraction Main Window – Noise Reducer Mode

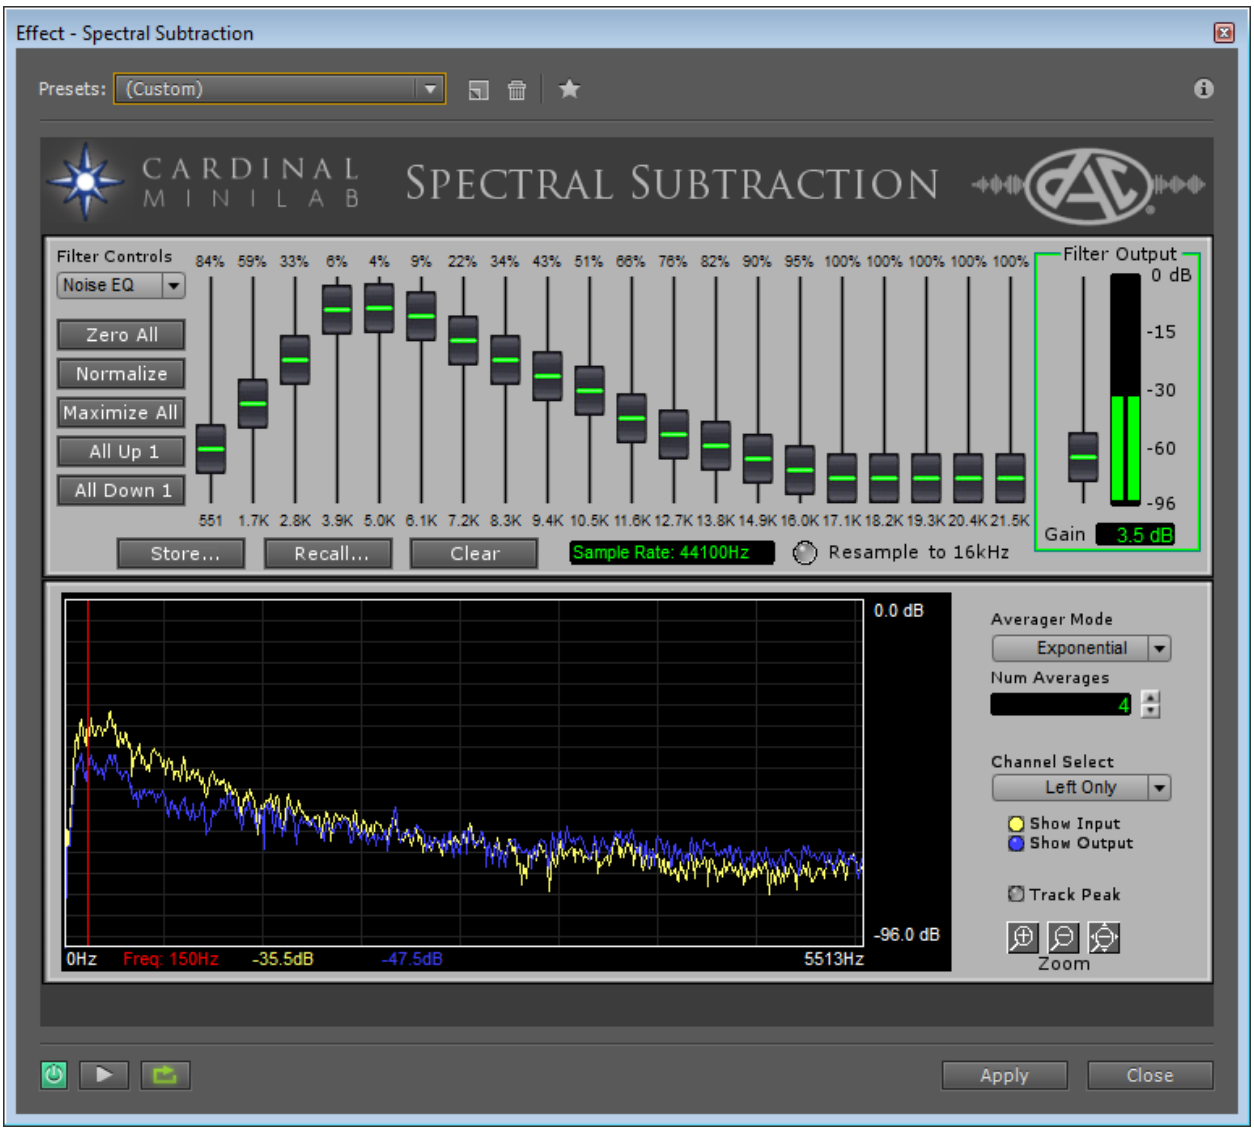

Figure 65: Spectral Subtraction Main Window - Noise EQ Mode

9.2 MAIN PLUG-IN WINDOW

Figure 64 and Figure 65 above show the plug-in as it would be displayed if used in Adobe Audition in **Noise Reducer** mode and **Noise EQ** mode respectively. The buttons above and below the plug-in window will appear differently depending upon your audio editing environment. Please consult the documentation on your specific audio editing environment for details of how to use these features.

This window provides access to all the functionality of the plug-in. A spectrum analyzer is provided as an aid in determining the characteristics of the target audio. Details of the plug-in are described in the following sections.

9.3 FILTER CONTROLS

9.3.1 Filter Controls

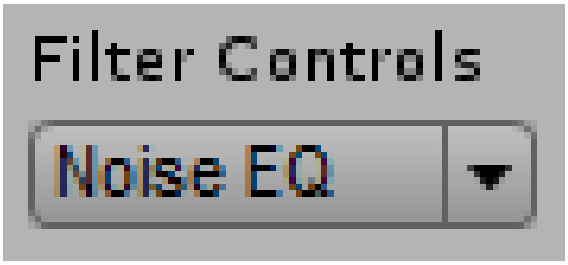

Figure 66: Filter Controls

The **Filter Controls** combo box allows the user to select either the **Noise EQ** mode or the **Noise Reducer** mode (displayed as N. Reducer). When switching from one mode to the other, the previous mode's controls are remembered and preserved.

9.3.2 Clear

Figure 67: Clear Button

Used to clear the spectral subtraction solution currently in memory and restart the algorithm from scratch.

9.3.3 Filter Output Gain

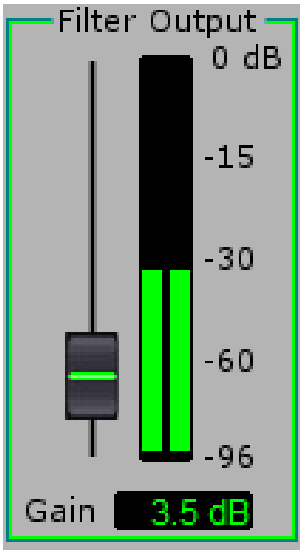

Figure 68: Filter Output Controls

The **Output Gain** control allows user to apply between 0 and 30dB of makeup gain to the processed output signal to maximize the signal level prior to final equalization, AGC, and listening/recording. The associated bargraph shows the actual output signal level after the gain has been applied.

9.3.4 Noise EQ Controls

The following sections apply only to the **Noise EQ** mode.

9.3.4.1 Slider Controls

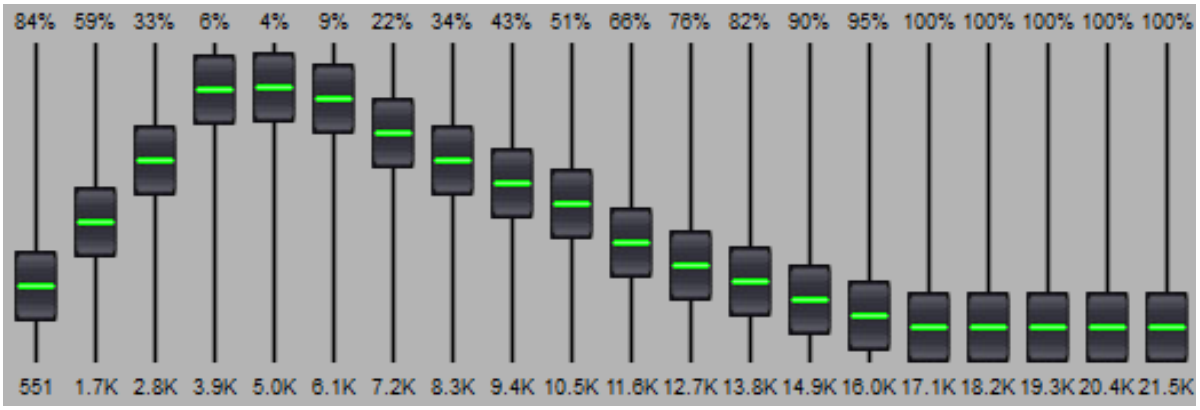

Figure 69: Noise EQ Slider Controls

The slider controls are used to specify the amount of noise reduction that the spectral subtraction attempts to apply to the input signal within each of 20 separate groups of frequency bands. Within each band, adjustment range is 0 (no attenuation) to 100% (maximal attenuation) in 1% increments.

9.3.4.2 Zero All

Figure 70: Zero All Button

This button instantly moves the slider controls for all bands to 0%, defeating the entire equalizer. This is a useful feature when it is desired to reset all sliders from scratch.

9.3.4.3 Normalize

Figure 71: Normalize Button

This button instantly shifts all slider controls up together until the top slider is at 0%. After normalization, the relative positioning of the sliders remains the same. This allows the filter to implement the desired equalization curve with minimum signal loss.

9.3.4.4 Maximize All

Figure 72: Maximize All Button

This button instantly moves the slider controls for all bands to 100%, maximizing the attenuation for all bands. This is a useful feature when it is desired to quickly adjust the sliders such that only a few bands are passed with all others rejected.

9.3.4.5 All Up 1

Figure 73: All Up 1 Button

This button shifts all sliders up 1%, while maintaining the desired curve. This is useful when the desired shape is found, but you want the filter to be less aggressive as a whole. No slider will be allowed to go less than 0%.

9.3.4.6 All Down 1

Figure 74: All Down 1 Button

This button shifts all sliders down 1%, while maintaining the desired curve. This is useful when the desired shape is found, but you want the filter to be more aggressive as a whole. No slider will be allowed to go more than 100%.

9.3.4.7 Store / Recall

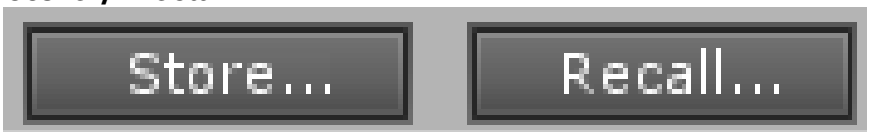

Figure 75: Store and Recall Buttons

These buttons allow the user to store and recall a slider configuration to a user-specified disk file that will not be lost when the computer is turned off.

9.3.5 Noise Reducer Controls

The following sections apply only to the **Noise Reducer** mode.

9.3.5.1 Slider Control

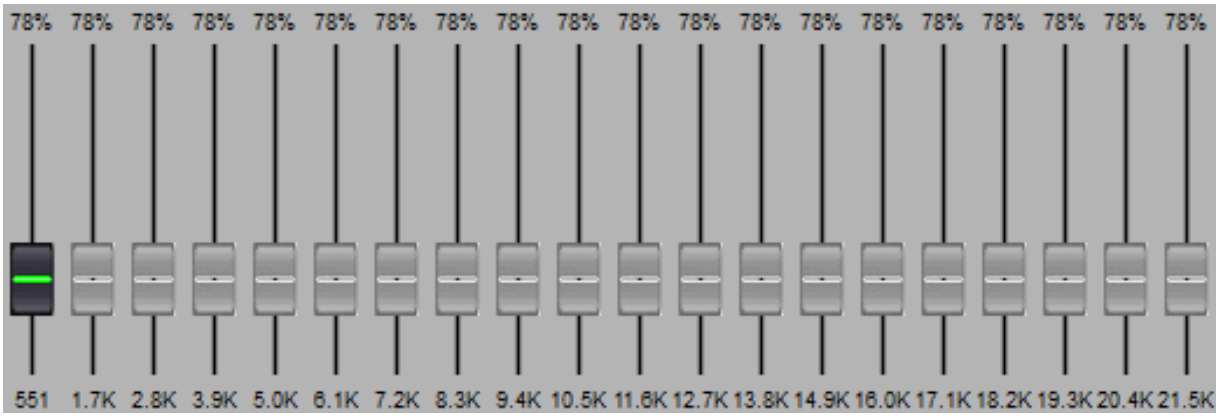

Figure 76: Noise Reducer Slider Controls

When in **Noise Reducer** mode, all the slider controls are linked to the first slider control. This slider is used to specify the amount of noise reduction that the spectral subtraction attempts to apply to the input signal across all frequency bands. Adjustment range is 0 (no attenuation) to 100% (maximal attenuation) in 1% increments.

10 MULTI-BAND FILTER

10.1 OVERVIEW OF THE MULTI-BAND FILTER

The Multi-Band filter provides two different filters in a single plug-in – the Multiple Notch and Multiple Slot.

Only one filter can be operational at a time.

10.1.1 Overview of Multiple-Notch

The Multiple Notch filter is used to remove, or "notch out," single-frequency noises such as tones or whistles with minimal effect on signal frequencies other than the notch frequency. Single notches can be added one at a time and configured individually. Also, notch "groups" can be added to cancel many harmonically related frequencies at once.

The Multiple Notch filter is synthesized from a frequency-domain representation of the desired notch profile. An inverse FFT builds FIR coefficients based on the frequency-domain representation. For this reason, the notches in this filter are "square" notches rather than "V" notches. Square notches mean that frequencies very close to the specified center frequency will be cancelled along with the center frequency. However, the square notches also mean that the Multiple Notch filter is able to tolerate moderate variances in the specified frequency such as those caused by "wow and flutter" effects. (Filters that use "V" notches include the Notch filter, the Comb filter, and the Parametric Equalizer.)

To properly utilize the Multiple Notch filter, you will first need to identify the noise frequencies. The easiest way to do this is to use the included spectrum analyzer. Once the noise frequencies have been identified, add a notch for each frequency.

Notches are defined by three values: the notch frequency, the notch width, and the notch depth. The notch frequency is simply the frequency at which the notch should be centered. The notch width defines the desired width of the square notch in Hz, and the notch depth defines the desired depth in dB.

Often, tonal noises include not only the fundamental frequency, but also harmonic multiples of that frequency. Instead of requiring the addition of an individual notch for each harmonic, the Cardinal Multiple Notch filter allows the addition of Notch Groups to cancel harmonically related tones in a single action. A Notch Group is defined in relation to its Base Notch. The Base Notch is defined with a frequency, width, and depth just like a single notch. Frequency, width, and depth of all other notches in the group will be calculated

based on these parameters.

10.1.2 Overview of Multiple-Slot

The Multiple Slot filter is used to isolate, or "slot" single-frequency noises such as tones or whistles in the input audio, attenuating all other audio. This is the exact opposite of the Multiple Notch filter function. Single slots can be added one at a time and configured individually. Also, slot “groups” can be added to isolate many harmonically related frequencies at once.

The Multiple Slot filter is synthesized from a frequency-domain representation of the desired slot profile. An inverse FFT builds FIR coefficients based on the frequency-domain representation. For this reason, the slots in this filter are “square” slots rather than “V” slots. Square slots mean that frequencies very close to the specified center frequency will be cancelled along with the center frequency. However, the square slots also mean that the Multiple Slot filter is able to tolerate moderate variances in the specified frequency such as those caused by “wow and flutter” effects.

To properly utilize the Multiple Slot filter, you will first need to identify the noise frequencies. The easiest way to do this is to use the included spectrum analyzer. Once the noise frequencies have been identified, add a slot for each frequency.

Slots are defined by three values: the slot frequency, the slot width, and the slot gain. The slot frequency is simply the frequency at which the slot should be centered. The slot width defines the desired width of the square slot in Hz, and the slot gain defines the desired amplitude in dB.

Often, tonal noises include not only the fundamental frequency, but also harmonic multiples of that frequency. Instead of requiring the addition of an individual slot for each harmonic, the Cardinal Multiple Slot filter allows the addition of Slot Groups to cancel harmonically related tones in a single action. A Slot Group is defined in relation to its Base Slot. The Base Slot is defined with a frequency, width, and gain just like a single slot. Frequency, width, and gain of all other slots in the group will be calculated based on these parameters.

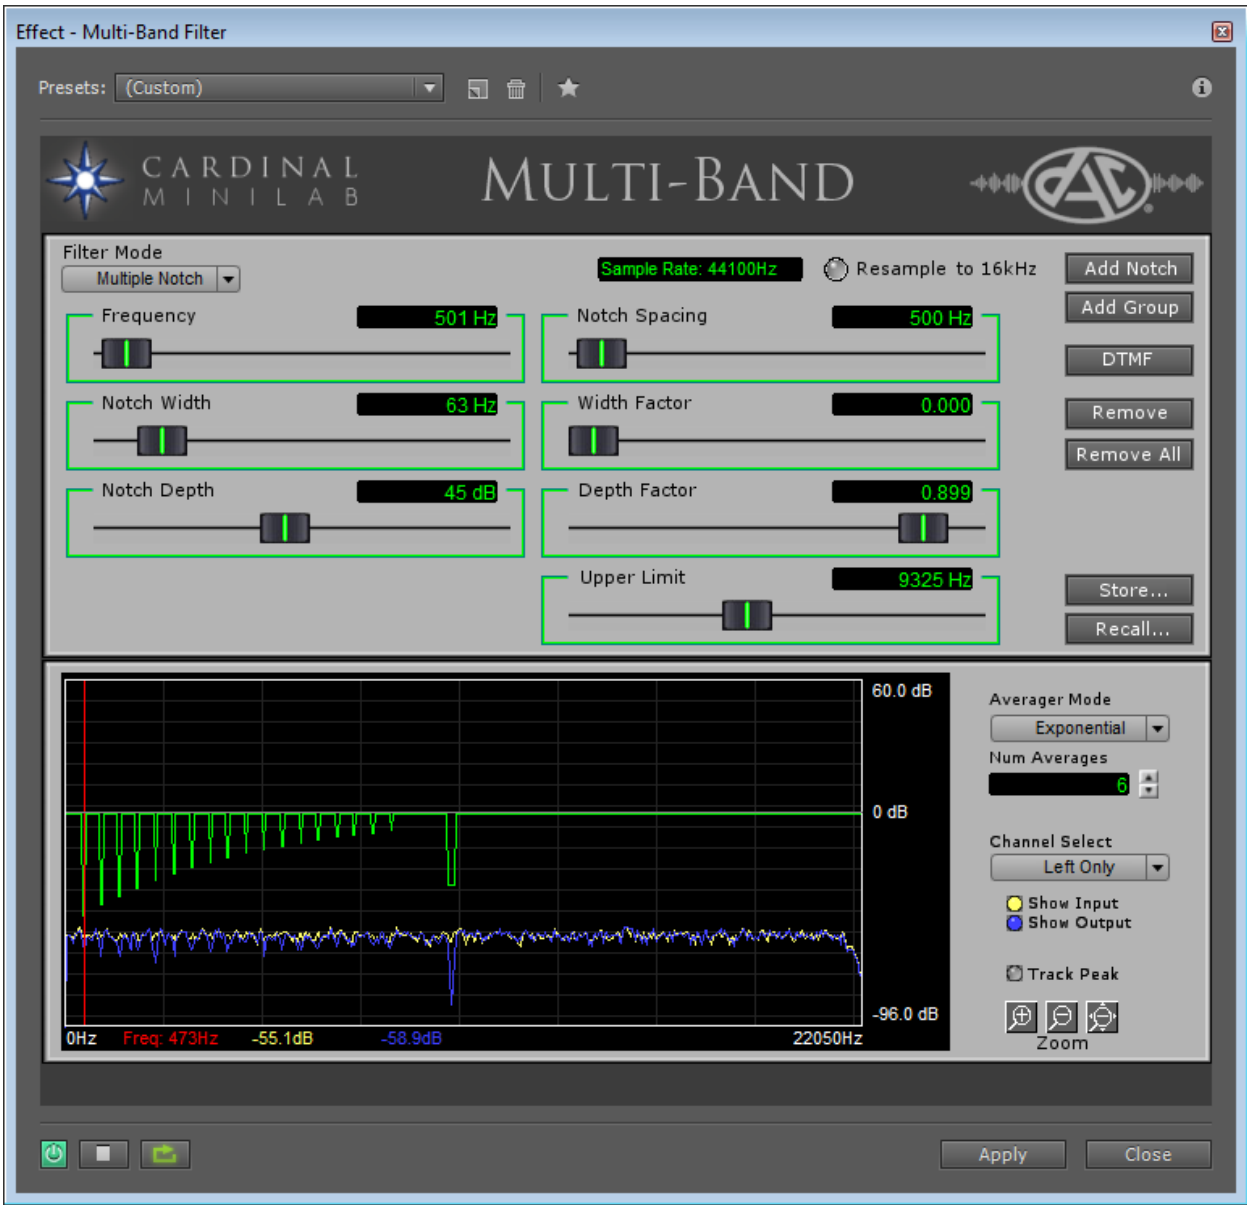

Figure 77: Multi-Band Main Window – Multiple Notch Mode

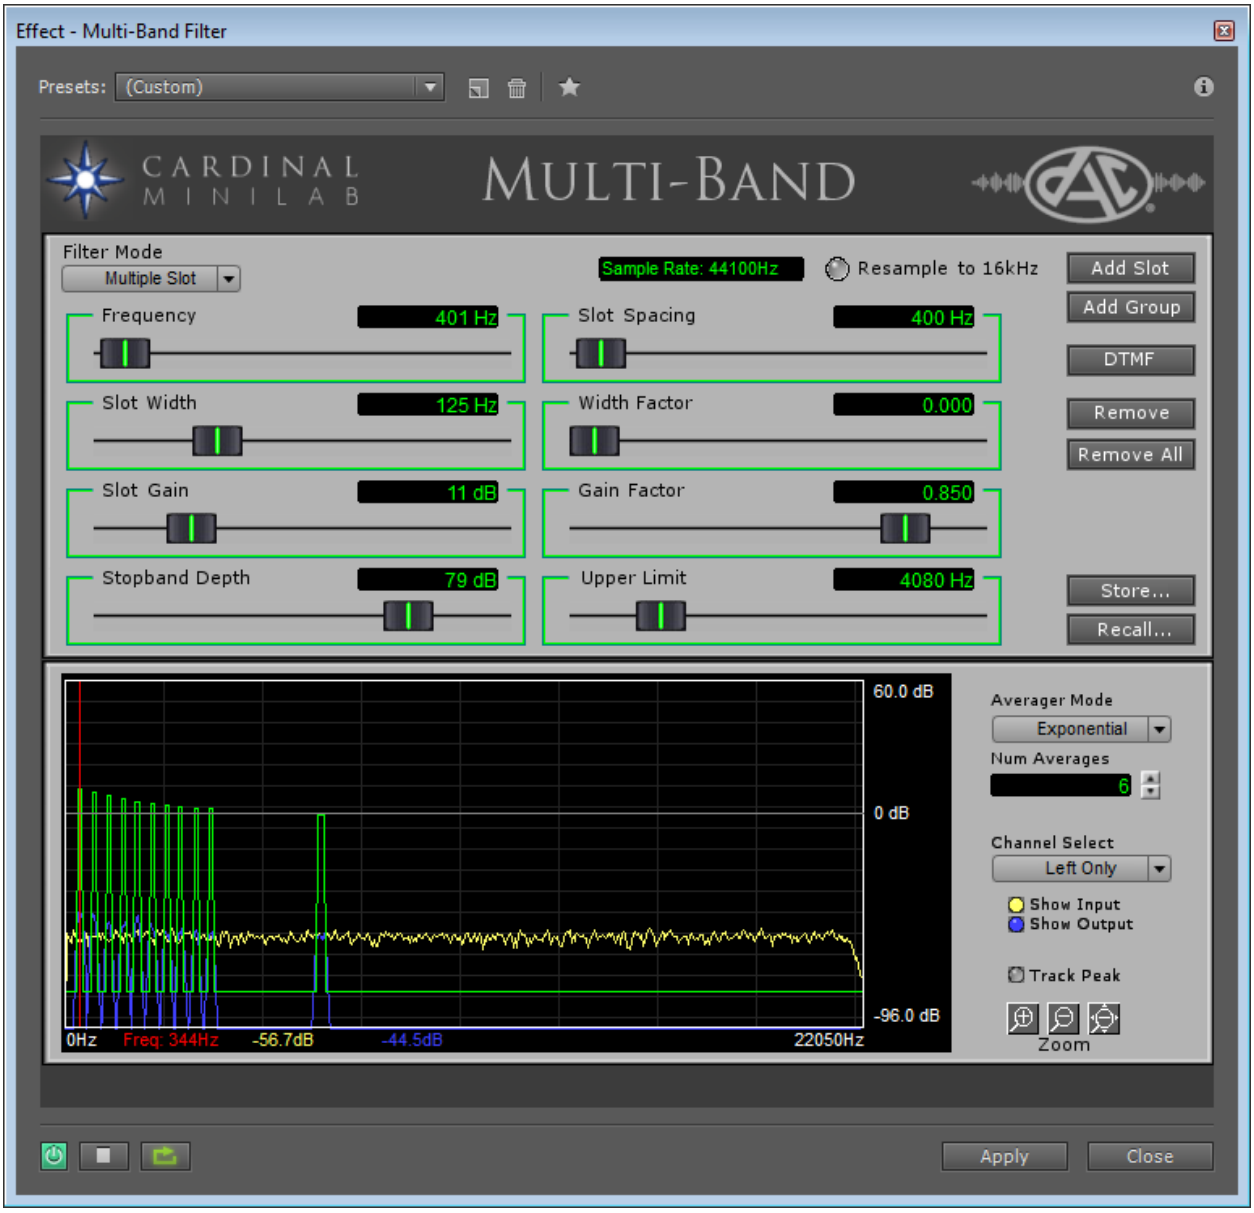

Figure 78: Multi-Band Main Window – Multiple Slot Mode

10.2 MAIN PLUG-IN WINDOW

Figure 77 and Figure 78 above show the plug-in as it would be displayed if used in Adobe Audition in **Multiple Notch** mode and **Multiple Slot** mode respectively. The buttons above and below the plug-in window will appear differently depending upon your audio editing environment. Please consult the documentation on your specific audio editing environment for details of how to use these features.

This window provides access to all the functionality of the plug-in. A spectrum analyzer is provided as an aid in determining the characteristics of the target audio. Details of the plug-in are described in the following sections.

10.3 FILTER CONTROLS

10.3.1 Filter Mode

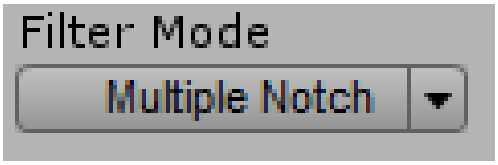

Figure 79: Filter Mode

The **Filter Mode** combo box allows the user to select either the **Multiple Notch** mode or the **Multiple Slot** mode. When switching from one mode to the other, the previous mode's controls are remembered and preserved.

10.3.2 Multiple Notch Filter Controls

10.3.2.1 Notch Frequency

Figure 80: Notch Frequency

Specifies frequency in Hertz which is to be removed from the input audio. Minimum Notch Frequency is 10 Hz, while maximum Notch Frequency depends upon the Sample Rate. Notch Frequency is adjustable in 1 Hz steps. Also sets the Base Frequency for a Notch Group.

10.3.2.2 Notch Width

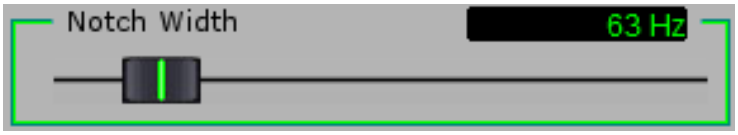

Figure 81: Notch Width

Width of the generated notch in Hertz.

10.3.2.3 Notch Depth

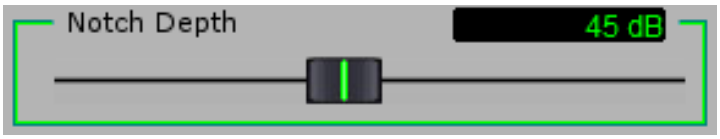

Figure 82: Notch Depth

Depth of the generated notch in dB.

10.3.2.4 Notch Spacing

Figure 83: Notch Spacing

Notch Spacing defines where the other notches in the group are to be placed. if the Base Notch frequency is F, and the spacing is set to S, then notches will be placed at frequencies F, F+S, F+2S, F+3S, etc.

10.3.2.5 Width Factor

Figure 84: Width Factor

Width Factor defines how wide the group notches should be. Frequency variations often occur as a percentage of the frequency, so the variation width in Hz is much larger at high frequencies. The Width Factor defines a percentage width up to a maximum of 1.9%, and each notch will be at least the width defined by that percentage. For instance, if a notch group has width factor = 0.015, and one of the notches in that group is at 1000 Hz, then the width of the 1000 Hz notch will be at least 1000 × 0.015 = 15 Hz.

NOTE: The frequency-domain representation used to build the Multiple Notch filter has an inherent minimum notch width. Especially at the lower frequency notches, the width specified by the Width Factor will often fall below that minimum width, in which case the minimum width is used. For this reason, the effect of the Width Factor control may only be visible at the higher frequency notches.

10.3.2.6 Depth Factor

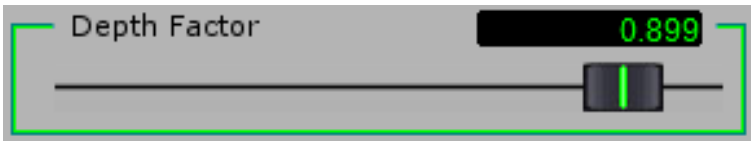

Figure 85: Depth Factor

Depth Factor defines how deep the group notches should be. Many harmonic tonal noises have a "1/f" volume profile, where the lower harmonics are strong and higher harmonics are progressively weaker. The Depth Factor controls the depth taper of the notches so that the notch depth can parallel the harmonic strength profile. The base notch always has the specified Notch Depth, while subsequent notches taper to smaller depths as frequency increases. The higher the Depth Factor, the more gradual the taper. A Depth Factor of 0.0 produces the most severe taper and means effectively that there are no harmonics at all. A Depth Factor of 1.0 means that notches have uniform depth at the Base Notch depth setting.

10.3.2.7 Upper Limit

Figure 86: Upper Limit

Upper Limit defines how many notches there are in the group. If a harmonic tonal noise only extends up to a certain frequency, it may be undesirable to notch out all multiples of the base frequency when only a few are needed. In this case, set the Upper Limit just above the highest frequency where a notch is desired; notches will be added up to that limit, and no notches will be added above the limit.

10.3.2.8 Add Notch

Figure 87: Add Notch Button

Adds a new single notch at the frequency indicated in the Notch Frequency box and by a marker on the visualization axis. The notch is added with default settings, and the user is presented with controls to adjust the frequency, width, and depth of the notch.

10.3.2.9 Add Group

Figure 88: Add Group Button

Adds a new notch group with its base notch at the frequency indicated in the Notch Frequency box and by a marker on the visualization axis. The notch group is added with default settings, and the user is presented with controls to adjust the frequency, width, depth, notch spacing, depth factor, width factor, and upper limit of the notch group.

10.3.2.10 DTMF Button

Figure 89: DTMF Button

Inserts 8 pre-defined frequencies that make up the **Dual-tone multi-frequency** (**DTMF**). The version of DTMF used for telephone **tone dialing** is known by the trademarked term **Touch-Tone**, and is standardized by ITU-T Recommendation.

10.3.2.11 Remove Button

Figure 90: Remove Button

Removes the currently selected notch or notch group from the filter.

10.3.2.12 Remove All Button

Figure 91: Remove All Button

Removes all notches and notch groups from the filter.

10.3.2.13 Store Button

Figure 92: Store Button

Saves the filter's current configuration to a disk file.

10.3.2.14 Recall Button

Figure 93: Recall Button

Loads a previously saved filter configuration from a disk file.

10.3.3 Multiple Slot Controls

10.3.3.1 Slot Frequency

Figure 94: Slot Frequency

Specifies frequency in Hertz which is to be enhanced in the input audio. Minimum Slot Frequency is 10 Hz, while maximum Slot Frequency depends upon the Sample Rate. Slot Frequency is adjustable in 1 Hz steps.

10.3.3.2 Slot Width

Figure 95: Slot Width

Width of the generated slot in Hz.

10.3.3.3 Slot Gain

Figure 96: Slot Gain

The gain of the base slot, in dB.

10.3.3.4 Stopband Depth

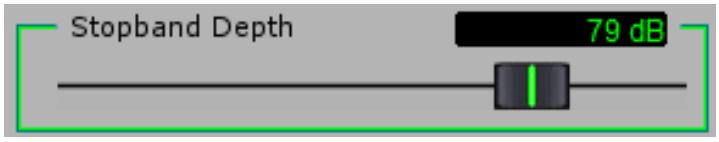

Figure 97: Stopband Depth

Specifies amount in dB by which frequencies other than the Slot Frequency are attenuated.

10.3.3.5 Slot Spacing

Figure 98: Slot Spacing

Slot Spacing defines where the other slots in the group are to be placed. if the Base Slot frequency is F, and the spacing is set to S, then slots will be placed at frequencies F, F+S, F+2S, F+3S, etc.

10.3.3.6 Width Factor

Figure 99: Width Factor

Width Factor defines how wide the group slots should be. Frequency variations often occur as a percentage of the frequency, so the variation width in Hz is much larger at high frequencies. The Width Factor defines a percentage width up to a maximum of 1.9%, and each notch will be at least the width defined by that percentage. For instance, if a slot group has width factor = 0.015, and one of the slots in that group is at 1000 Hz, then the width of the 1000 Hz slot will be at least 1000 × 0.015 = 15 Hz.

NOTE: The frequency-domain representation used to build the Multiple Slot filter has an inherent minimum slot width. Especially at the lower frequency slots, the width specified by the Width Factor will often fall below that minimum width, in which case the minimum width is used. For this reason, the effect of the Width Factor control may only be visible at the higher frequency notches.

10.3.3.7 Gain Factor

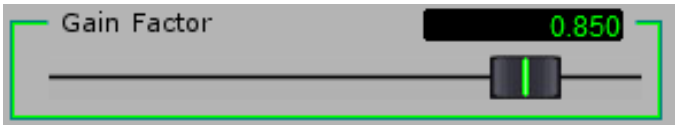

Figure 100: Gain Factor

Gain Factor defines how much gain is applied to the group slots. Many harmonic tonal noises have a "1/f" volume profile, where the lower harmonics are strong and higher harmonics are progressively weaker. The Gain Factor controls the gain taper of the slots so that the slot gain can parallel the harmonic strength profile. The base slot always has the specified Slot Gain, while subsequent slots taper to smaller gains as frequency increases. The higher the Gain Factor, the more gradual the taper. A Gain Factor of 0.0 produces the most severe taper and means effectively that there are no harmonics at all. A Gain Factor of 1.0 means that slots have uniform gain at the Base Slot gain setting.

10.3.3.8 Upper Limit

Figure 101: Upper Limit

Upper Limit defines how many slots there are in the group. If a harmonic tonal noise only extends up to a certain frequency, it may be undesirable to slot out all multiples of the base frequency when only a few are needed. In this case, set the Upper Limit just above the highest frequency where a slot is desired; slots will be added up to that limit, and no slots will be added above the limit.

10.3.3.9 Add Slot

Figure 102: Add Slot Button

Adds a new single slot at the frequency indicated in the Slot Frequency box and by a marker on the visualization axis. The slot is added with default settings, and the user is presented with controls to adjust the frequency, width, and gain of the slot.

10.3.3.10 Add Group

Figure 103: Add Group Button

Adds a new slot group with its base slot at the frequency indicated in the Slot Frequency box and by a marker on the visualization axis. The slot group is added with default settings, and the user is presented with controls to adjust the frequency, width, gain, slot spacing, gain factor, width factor, and upper limit of the slot group.

10.3.3.11 DTMF Button

Figure 104: DTMF Button

Inserts 8 pre-defined frequencies that make up the **Dual-tone multi-frequency** (**DTMF**). The version of DTMF used for telephone **tone dialing** is known by the trademarked term **Touch-Tone**, and is standardized by ITU-T Recommendation.

10.3.3.12 Remove Button

Figure 105: Remove Button

Removes the currently selected slot or slot group from the filter.

10.3.3.13 Remove All Button

Figure 106: Remove All Button

Removes all slots and slot groups from the filter.

10.3.3.14 Store Button

Figure 107: Store Button

Saves the filter's current configuration to a disk file.

10.3.3.15 Recall Button

Figure 108: Recall Button

Loads a previously saved filter configuration from a disk file.

11 GRAPHIC EQ

11.1 OVERVIEW OF THE GRAPHIC EQUALIZER

The Graphic EQ provides two different equalizers in a single plug-in – the **Graphic EQ** and **Hi-Res Graphic**.

Only one equalizer can be operational at a time.

11.1.1 Overview of Graphic EQ

The 20-band Graphic Equalizer is an easy-to-use linear-phase FIR digital filter that is used to reshape the spectrum of the final output signal. Reshaping is accomplished with twenty vertical scroll bars (also called "slider" controls) which adjust the attenuation of each frequency band. These controls are very similar to the slider controls found on analog graphic equalizers found on many consumer stereo systems, and thus should be very familiar to even the novice user.

However, unlike analog graphic equalizers, this digital equalizer has some very powerful additional capabilities. For example, the **Normalize** button allows the user to instantly move all slider controls up until the top slider is at 0dB. The **Zero All** button instantly sets all the sliders to 0dB, while the **Maximize** button instantly sets all the sliders to -100dB. The **All Down 1dB** button instantly moves all sliders down in 1dB increments, while the **All Up 1dB** button moves all sliders up in 1dB increments. None of these functions are available in an analog graphic equalizer! Notice also that the 20 sliders are spread across the selected Bandwidth and that the frequency spacing is optimized for voice processing.

Additionally, since a computer with a disk drive operates the equalizer, a **Store** and **Recall** capability is available. This allows the user to store commonly-used slider configurations in disk memories so that they can be instantly recalled later whenever they are needed, without having to manually adjust the slider controls.

11.1.2 Overview of Hi-Res Graphic

In some applications, it may be necessary to precisely reshape the spectrum of input audio prior to passing it through successive filter stages. For example, if the audio is from a microphone which has an unusual frequency response curve (for example, a microphone acoustically modified as a result of concealment), a compensation filter that reshapes the audio to a normal spectral shape might be desirable.

The Hi-Res Graphic Filter is essentially a 460-band graphic equalizer; however, instead of having 460 separate slider controls, it allows the user to precisely

draw the desired filter shape on the computer screen, using the mouse, with as much or as little detail as desired. Once the filter shape has been drawn, a linear-phase digital filter is constructed in the PC and transferred to the external processor.

The **Normalize** button allows the user to shift the entire filter curve up until the highest point is at 0dB.

A **Store** and **Recall** capability is also provided to allow the user to store commonly-used filter shapes to disk memories so that they can be recalled later.

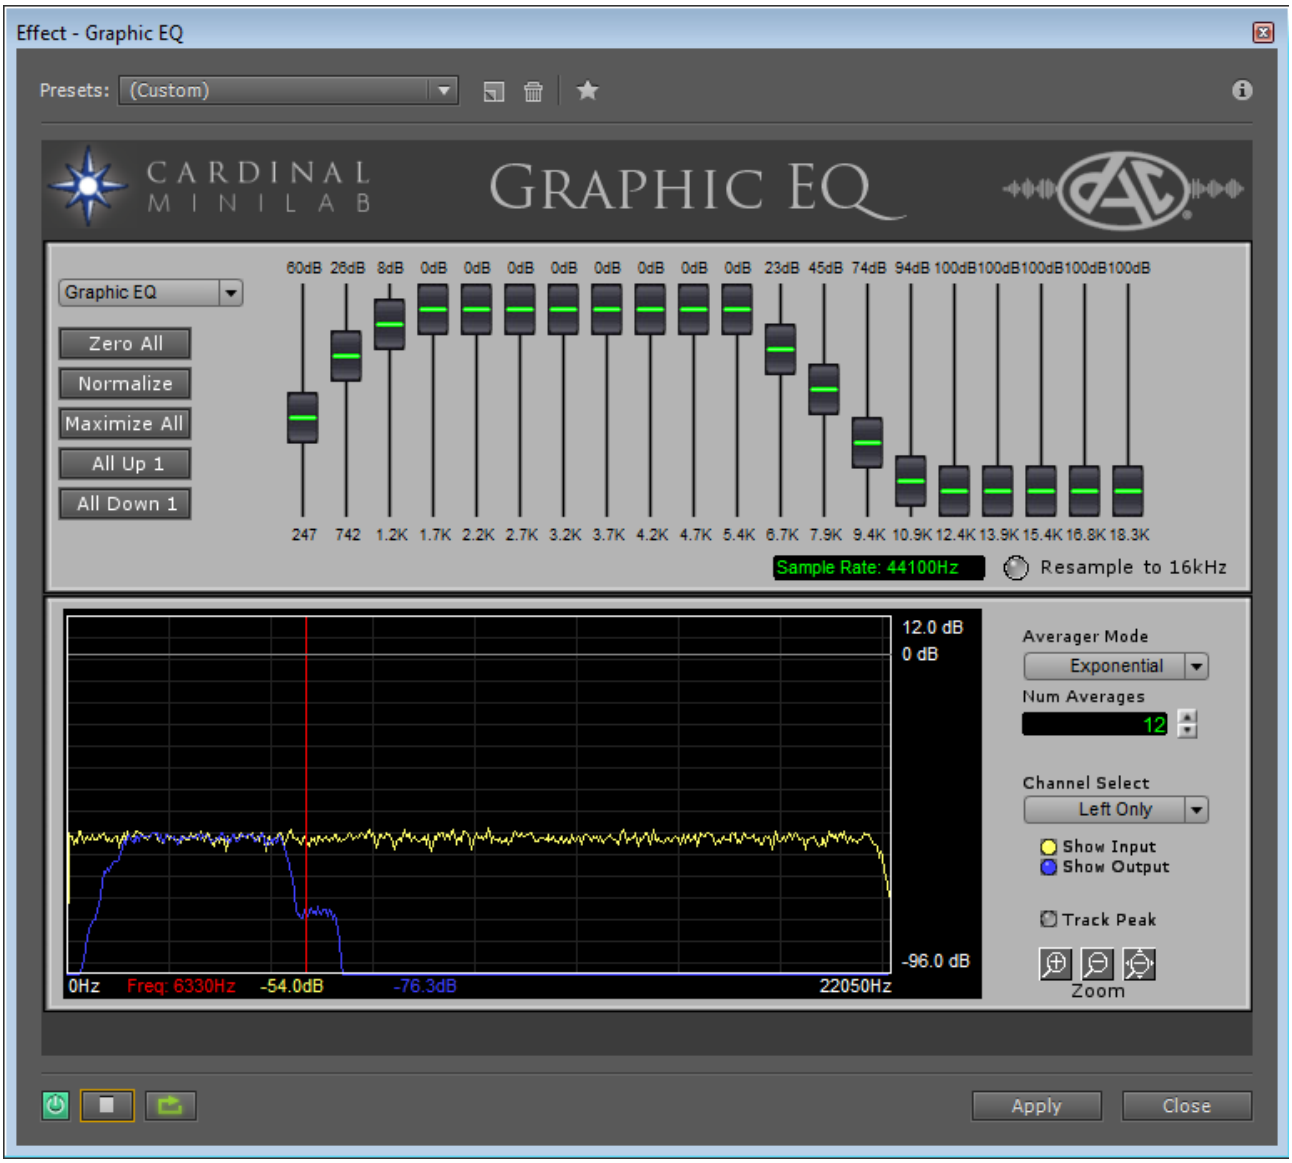

Figure 109: Graphic EQ Main Window – Graphic EQ Mode

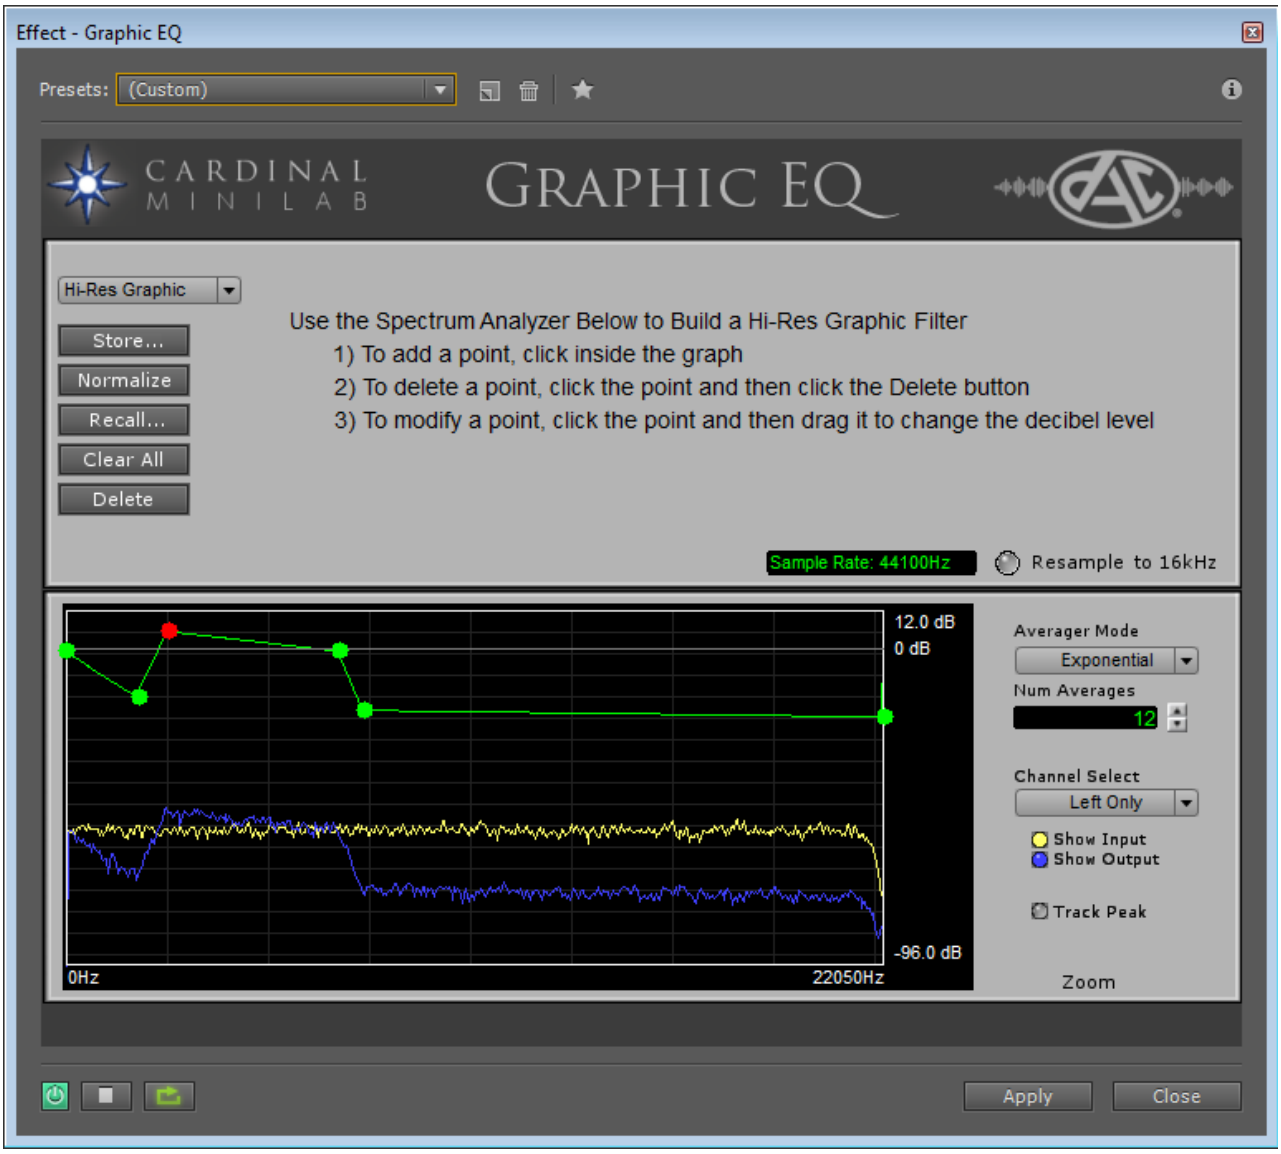

Figure 110: Graphic EQ Main Window – Hi-Res Graphic Mode

11.1 MAIN PLUG-IN WINDOW

Figure 109 and Figure 110 above show the plug-in as it would be displayed if used in Adobe Audition in **Graphic EQ** mode and **Hi-Res Graphic** mode respectively. The buttons above and below the plug-in window will appear differently depending upon your audio editing environment. Please consult the documentation on your specific audio editing environment for details of how to use these features.

This window provides access to all the functionality of the plug-in. A spectrum

analyzer is provided as an aid in determining the characteristics of the target audio. Details of the plug-in are described in the following sections.

11.2 EQUALIZER CONTROLS

11.2.1 Equalizer Mode

Figure 111: Equalizer Mode

The **Equalizer Mode** combo box allows the user to select either the **Graphic EQ** mode or the **Hi-Res Graphic** mode. When switching from one mode to the other, the previous mode's controls are remembered and preserved.

11.2.2 Graphic EQ Controls

11.2.2.1 Slider Controls

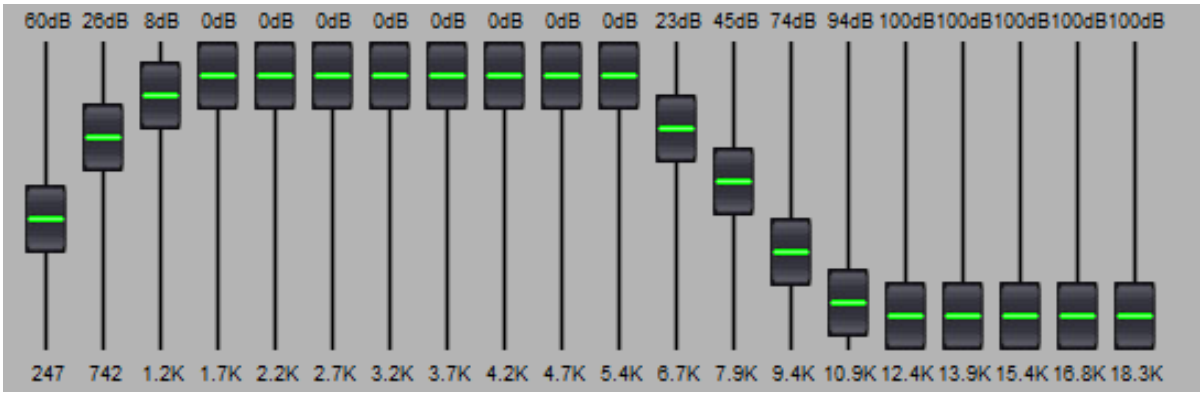

Figure 112: Slider Controls

The twenty vertical scroll bar "slider" controls are used to set the frequency response of the equalizer. Each slider can set the gain of its frequency band to any value between 0dB and -100 dB in 1dB steps.

11.2.2.2 Zero All

Figure 113: Zero All Button

This button instantly moves the slider controls for all bands to 0dB, defeating the entire equalizer. This is a useful feature when it is desired to reset all sliders from scratch.

11.2.2.3 Normalize Button

Figure 114: Normalize Button

This button instantly shifts all slider controls up together until the top slider is at 0dB. After normalization, the relative positioning of the sliders remains the same. This allows the digital equalizer to implement the desired equalization curve with minimum signal loss.

11.2.2.4 Maximize All

Figure 115: Maximize All Button

This button instantly moves the slider controls for all bands to -100dB, maximizing the attenuation for all bands. This is a useful feature when it is desired to quickly adjust the sliders such that only a few bands are passed with all others rejected.

11.2.2.5 All Up 1

Figure 116: All Up 1 Button

This button shifts all sliders up by 1dB from their current position; no slider, however, will be allowed to go above 0dB. This button allows the user to shift the entire equalizer curve up so that there will be room to move one or more sliders down relative to the others.

11.2.2.6 All Down 1

Figure 117: All Down 1 Button

This button shifts all sliders down by 1dB from their current position; no slider, however, will be allowed to go below -100dB. This button allows the user to shift the entire equalizer curve down so that there will be room to move one or more sliders up relative to the others.

11.2.3 Hi-Res Graphic Controls

11.2.3.1 Hi-Res Graphic Mini-Tutorial

The smoothing curve is graphed by the user using *control points*. These control

points are seen in Figure 110 as large circles on the graph. Control points represent a point on the curve where the slope of the line changes. Users can manipulate these control points in one of three ways:

- Add a control point
- Delete a control point
- Move a control point

To add a control point, simply click on the graph where you want it to be. The control point will immediately appear and you will hear the audio change immediately.

To delete a control point, click on an existing control point (it will turn red) and then click the **Delete** button (except the first and the last points, they cannot be deleted). This will remove the control point and the curve will snap back between the control points on either side.

To move a control point, click on an existing control point (it will turn red) and drag it with the mouse. Control points can only be moved vertically, which adjusts the gain at that point. Control points cannot be moved horizontally in an attempt to change the frequency at which the control point exists.

11.2.3.2 Store Button

Figure 118: Store Button

This button allows the user to store the current configuration to a user-specified disk file that will not be lost when the computer is turned off.

11.2.3.3 Recall Button

Figure 119: Recall Button

This button allows the user to recall a previously stored configuration from any of the saved disk files previously generated using the **Store** button.

11.2.3.4 Normalize Button

Figure 120: Normalize Button

The **Normalize** button allows the user to shift the entire filter curve up until the highest point is at 0dB.

11.2.3.5 Clear All

Figure 121: Clear All Button

The **Clear All** button will clear all graph control points and reset the filter to a passthrough.

11.2.3.6 Delete Button

Figure 122: Delete Button

The **Delete** button will delete the current control point (highlighted in red) in the graph.

12 GAIN STAGE

12.1 OVERVIEW OF THE GAIN STAGE

The Gain Stage provides two different gain stages in a single plug-in – the **Automatic Gain Control (AGC)** and **Limiter/Compressor/Expander (LCE)**.

Only one gain stage can be operational at a time.

12.1.1 Overview of AGC

The Automatic Gain Control automatically attempts to boost low-level output signals to a peak reference level (-18dB bargraph level) by gradually increasing output signal gain over a specified Release Time interval until either the proper level or Maximum Gain has been reached. This compensates for near party/far party conversations and for losses in signal level which may have occurred during the enhancement process. If the output signal levels are at or above the -12 dB reference level, the AGC will have no effect.

12.1.2 Overview of LCE

The Limiter/Compressor/Expander (LCE) is a three-section signal level processor allowing manipulation of the overall dynamic range of a signal. The LCE is typically used to correct for near-party/far-party or quiet talker scenarios.

The three sections correspond to three types of level processing available – limiting, compression, and expansion. Limiting is applied to the loudest levels in a signal. Compression is the middle region, and expansion is applied to the quietest levels.

- In the Limiting region, the output signal level is "damped" to the Limiting Threshold level. When the input signal level is in the Limiting region, attenuation is applied to keep the output level from exceeding the specified Limit Threshold.

- In the Compression region, levels are adjusted so that output signal level changes are smaller than their corresponding input signal level changes. Thus, the LCE decreases the dynamic range of the signal for levels in the Compression region. As an example, a 2:1 compressor would produce an output level change of only 10 dB when the input signal changes by 20 dB. Compression is often used to correct near-party/far-party level differences, boosting the lower-level far-party speech relative to the louder near-party speech. Compression also eases listening, especially for noisy audio. Compressors are generally preferred over AGCs since input signal level differences are more closely preserved.

- In the Expansion region, levels are adjusted so that output signal level changes are larger than their corresponding input signal level changes. Thus, the LCE increases the dynamic range of the signal for levels in the Expansion region. Expansion is the opposite of compression. For example, a 1:3 expander would produce an output level change of 30 dB when the input signal changes by 10 dB. A 1:2 expansion would restore a signal's dynamic range following a 2:1 compression. Expansion is also used to attenuate objectionable low-level background noise that is below the voice level.

Figure 123 shows an example LCE curve. In this example, the Limiting Threshold is set at -20dB, and the Compression Threshold is set at -60dB. The Compression Ratio is 2:1, and the Expansion Ratio is 1:3.

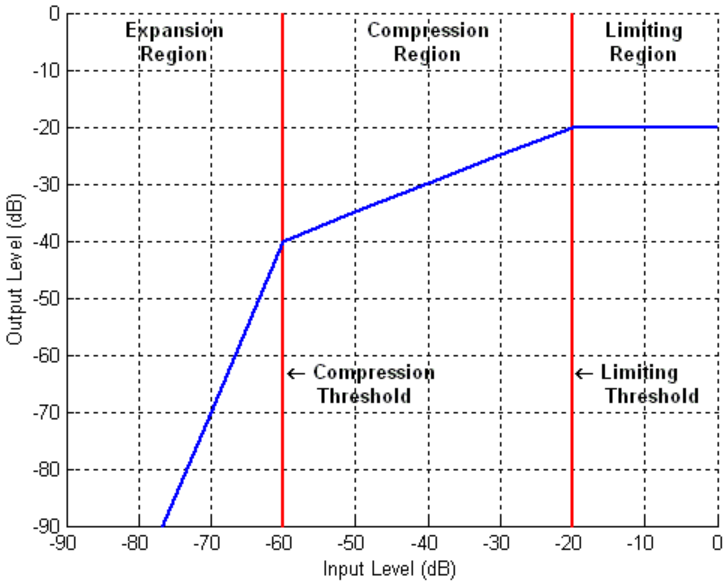

Figure 123: Example LCE Curve

In each section, the LCE modifies the amplitude of the signal using a variable-gain digital amplifier. The amplitude is a rectified and smoothed version of the signal waveform, as measured by a real-time digital envelope detector. So, in the figure above, the "Input Level" actually refers to the smoothed level envelope rather than the sample-by-sample instantaneous input level. The operation of the envelope detector is governed by the Attack Time, Release Time, and Lookahead controls.

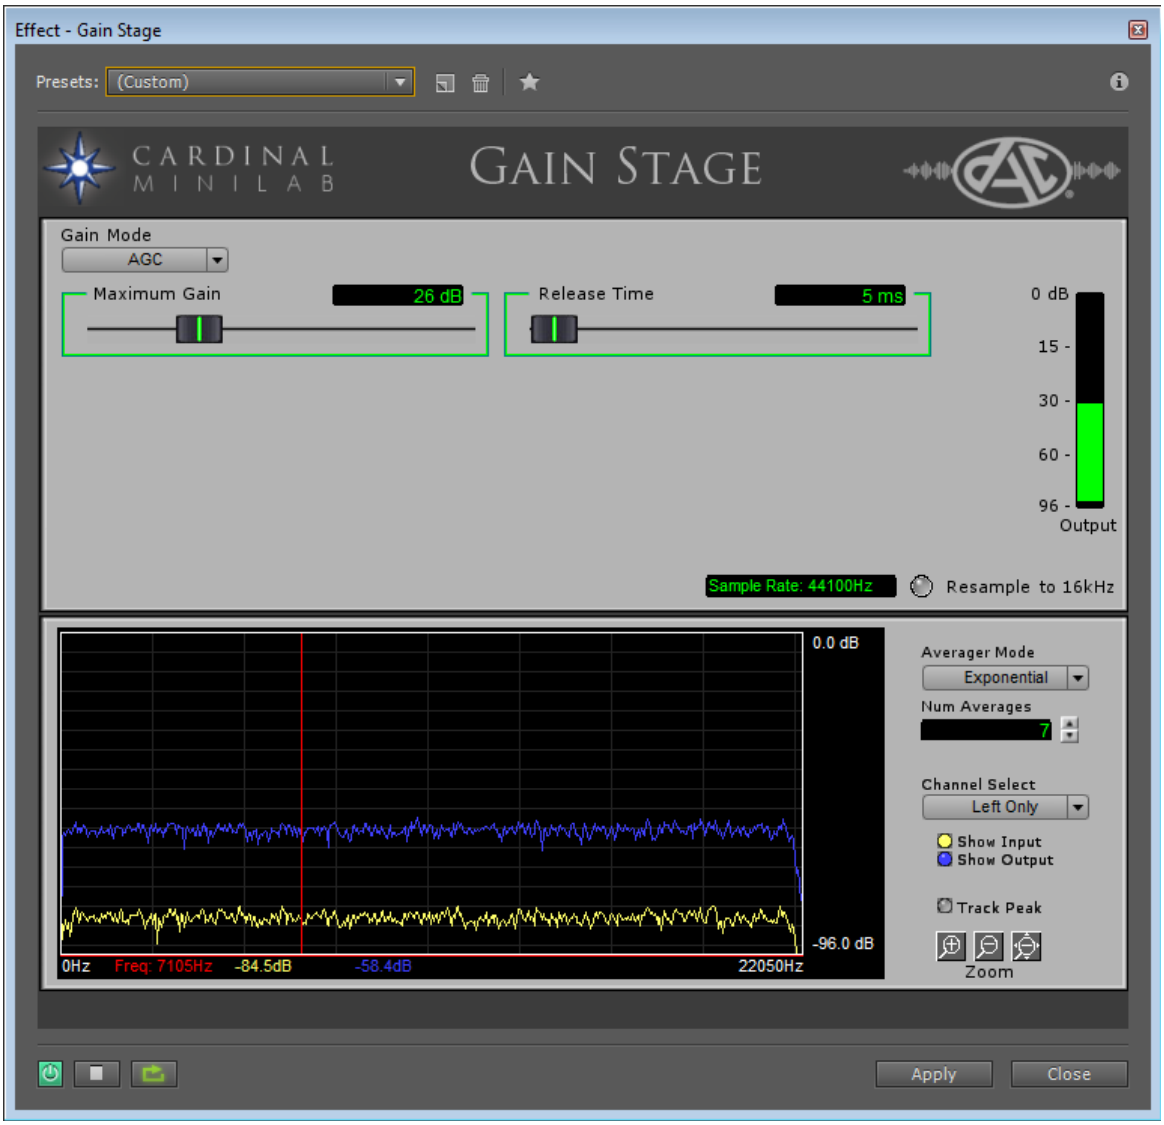

Figure 124: Gain Stage Main Window – AGC Mode

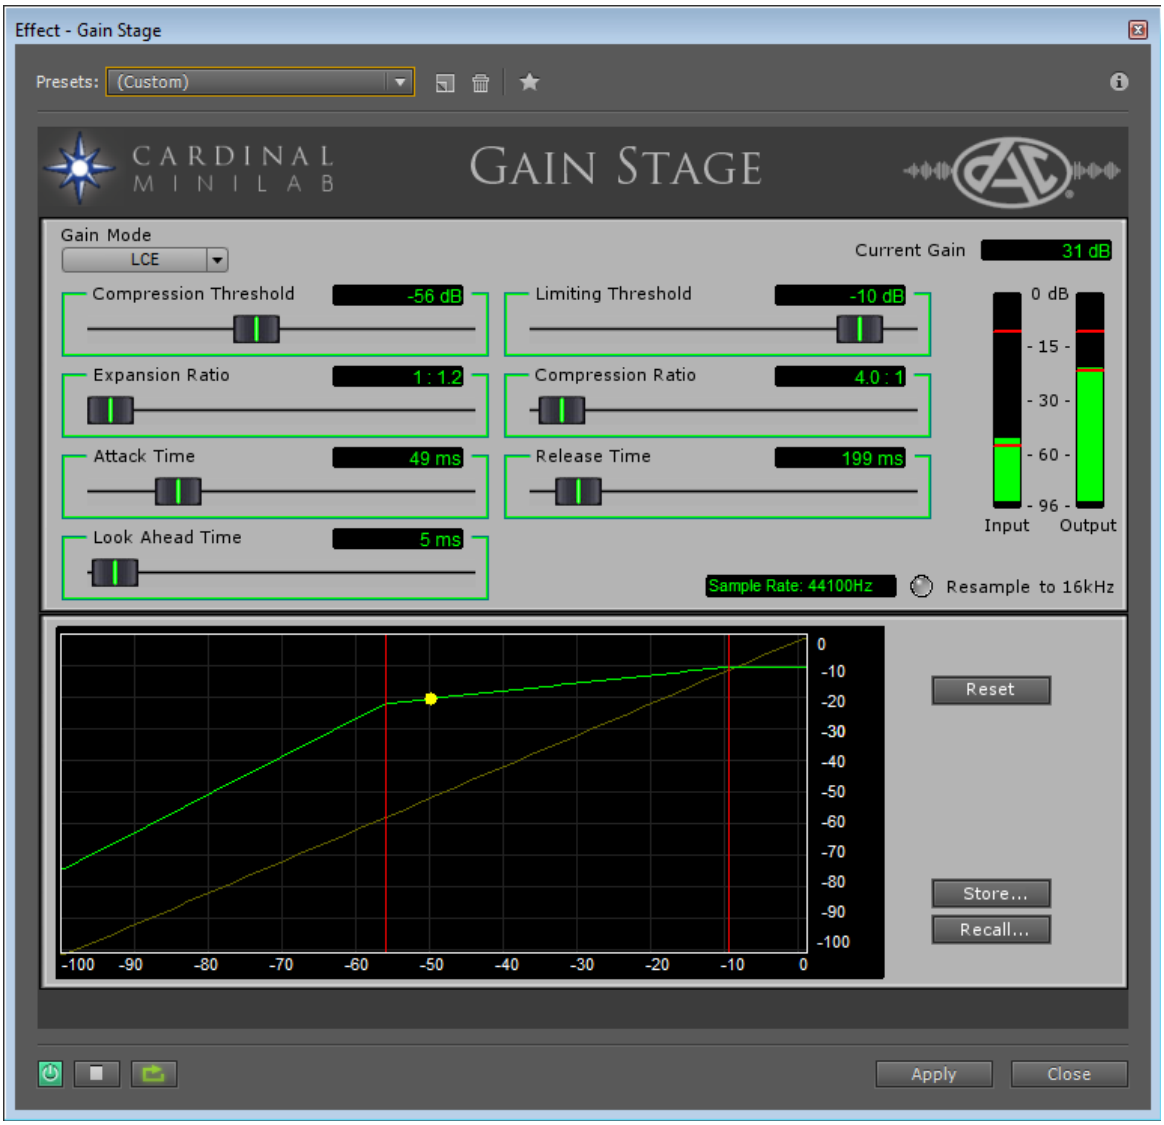

Figure 125: Gain Stage Main Window – LCE Mode

12.2 MAIN PLUG-IN WINDOW

Figure 124 and Figure 125 above show the plug-in as it would be displayed if used in Adobe Audition in **AGC** mode and **LCE** mode respectively. The buttons above and below the plug-in window will appear differently depending upon your audio editing environment. Please consult the documentation on your specific audio editing environment for details of how to use these features.

This window provides access to all the functionality of the plug-in. A spectrum analyzer is provided as an aid in determining the characteristics of the target audio. Details of the plug-in are described in the following sections.

12.3 GAIN CONTROLS

12.3.1 Gain Mode

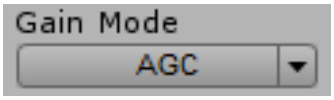

Figure 126: Gain Mode

The **Gain Mode** combo box allows the user to select either the **AGC** mode or the **LCE** mode. When switching from one mode to the other, the previous mode's controls are remembered and preserved.

12.3.2 AGC Controls

12.3.2.1 Maximum Gain

Figure 127: Maximum Gain

Maximum Gain specified how much gain the AGC can apply in its attempt to bring the output signal up to the desired level. The greater the Maximum Gain, the lower the signal that can be brought up to the threshold level. The Maximum Gain range is 0-100dB. For most near-party/far-party applications, around 10dB is recommended. Settings greater than 10dB may elevate background noise to an objectionable level during pauses in speech. A "soft AGC" using of 5dB is often useful even when large voice level differences are not present.

12.3.2.2 Release Time

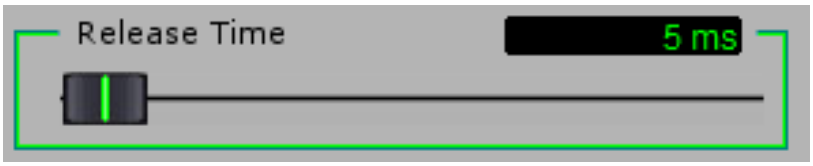

Figure 128: Release Time

Release Time controls how quickly the LCE will respond to decreases in input signal level. The shorter the Release Time, the more quickly the AGC will react. For most voice applications, a release time of about 200 milliseconds in recommended. Release Time settings less than 200 milliseconds may result in annoying "pumping" sounds as the AGC changes gain during rapid-fire

conversations.

12.3.3 LCE Controls

12.3.3.1 Compression Threshold

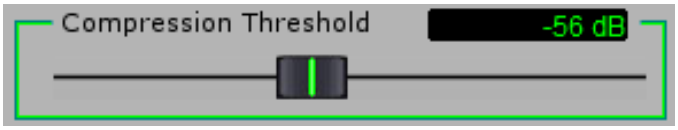

Figure 129: Compression Threshold

The level above which compression is applied to the signal. The specified compression ratio is applied to the input signal whenever the input level is between the Compression Threshold and the Limit Threshold.

12.3.3.2 Limiting Threshold

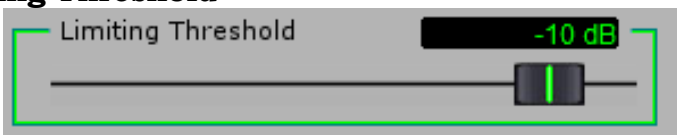

Figure 130: Limiting Threshold

The level above which the signal is damped. For instance, if the Limit Threshold is –20dB, all signal levels above –20dB will be attenuated to –20dB.

12.3.3.3 Expansion Ratio

Figure 131: Expansion Ratio

Specifies the amount of expansion to be applied to the signal when the input level falls in the Expansion Region. The Expansion Ratio is expressed as a ratio 1:N. Jumps in the output signal are N times larger than their corresponding jumps in the input signal. For example, with an Expansion Ratio of 1:3, a 10 dB jump in input level becomes a 30dB jump in output level.

12.3.3.4 Compression Ratio

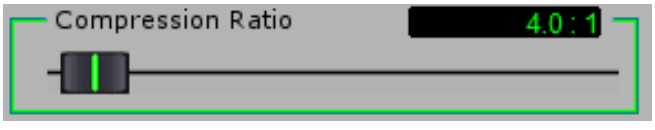

Figure 132: Compression Ratio

Specifies the amount of compression to be applied to the signal when the input level falls in the Compression Region. The Compression Ratio is expressed as a ratio N:1. Jumps in the output signal are N times smaller than their

corresponding jumps in the input signal. For example, with a Compression Ratio of 3:1, a 30dB jump in input level becomes a 10dB jump in output level.

12.3.3.5 Attack Time

Figure 133: Attack Time

Controls how quickly the LCE responds to increases in input signal level. For a more peak-sensitive processor, use a short Attack Time. For a more average-sensitive processor, use a longer Attack Time. For most speech applications, a fast Attack Time of 2-5 milliseconds is recommended.

12.3.3.6 Release Time

Figure 134: Release Time

Controls how quickly the LCE responds to decreases in input signal level. Short Release Times (<100 milliseconds) can create an annoying "pumping" artifact as the level detector is too responsive to intra-syllabic pauses. Long Release Times (>500 milliseconds) may fail to respond to breath group pauses and exchanges between speakers. For most speech applications, a Release Time of 200-400 milliseconds is recommended.

12.3.3.7 Look Ahead Time

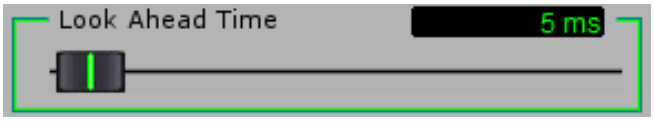

Figure 135: Look Ahead Time

Lookahead controls the alignment of the envelope detector with the output signal. Since the envelope is a smoothed version of the signal waveform, level changes in the envelope will lag corresponding changes in the signal itself. The applied LCE gain depends on the envelope level, so the same lag is reflected in the applied gain.

The Lookahead control adjusts an internal delay that compensates for this lag. The larger the Lookahead setting, the earlier the gain adjustments will be shifted. For most speech applications, a Lookahead of 1-5 milliseconds is recommended.

12.3.3.8 Reset Button

Figure 136: Reset Button

The **Reset** button resets all the LCE parameters to their default values, which are:

Parameter	Default Value
Limiting Threshold	-10dB
Compression Threshold	-50dB
Compression Ratio	1.5:1
Expansion Ratio	1:1.2
Attack Time	50ms
Release Time	200ms
Look Ahead Time	5ms

12.3.3.9 Store Button

Figure 137: Store Button

This button allows the user to store the current configuration to a user-specified disk file that will not be lost when the computer is turned off.

12.3.3.10 Recall Button

Figure 138: Recall Button

This button allows the user to recall a previously stored configuration from any of the saved disk files previously generated using the **Store** button.

13 PARAMETRIC EQUALIZER

13.1 OVERVIEW OF THE PARAMETRIC EQUALIZER

The Parametric Equalizer consists of a variable number of IIR filter stages, connected in series, which can be used for boosting (or peaking) and cutting (or nulling) portions of the input signal's frequency spectrum. Each stage is described by a center frequency, a frequency width, and a boost/cut amount, and can be configured independently. A common application of the parametric equalizer is to construct a precision notch filter which will perform nulling of the input signal at the specified center frequencies.

In the Current Stage block, the eight available stages can be selected one at a time to adjust their individual configurations. Individual stages can be toggled between Active and Inactive. An active stage is applied to the audio, while an inactive stage is bypassed. When a stage is made inactive, its settings are preserved.

HINT: It is often helpful to activate only one stage at a time when adjusting the stage settings. Then, once satisfactory settings have been found for each individual stage, all stages can be activated for audio processing.

When multiple stages are in use, their effects can overlap so that the overall signal level is reduced or boosted more than expected. For this reason, an output gain control is available as part of the Parametric Equalizer, allowing the user to compensate for overall level changes that may result from Parametric Equalizer filtering. (Advanced users may note that many Parametric EQ filters provide an input attenuation control so that fixed-point saturation can be avoided. Since Cardinal uses a floating-point implementation, saturation is not a concern, so only output level adjustment is provided.)

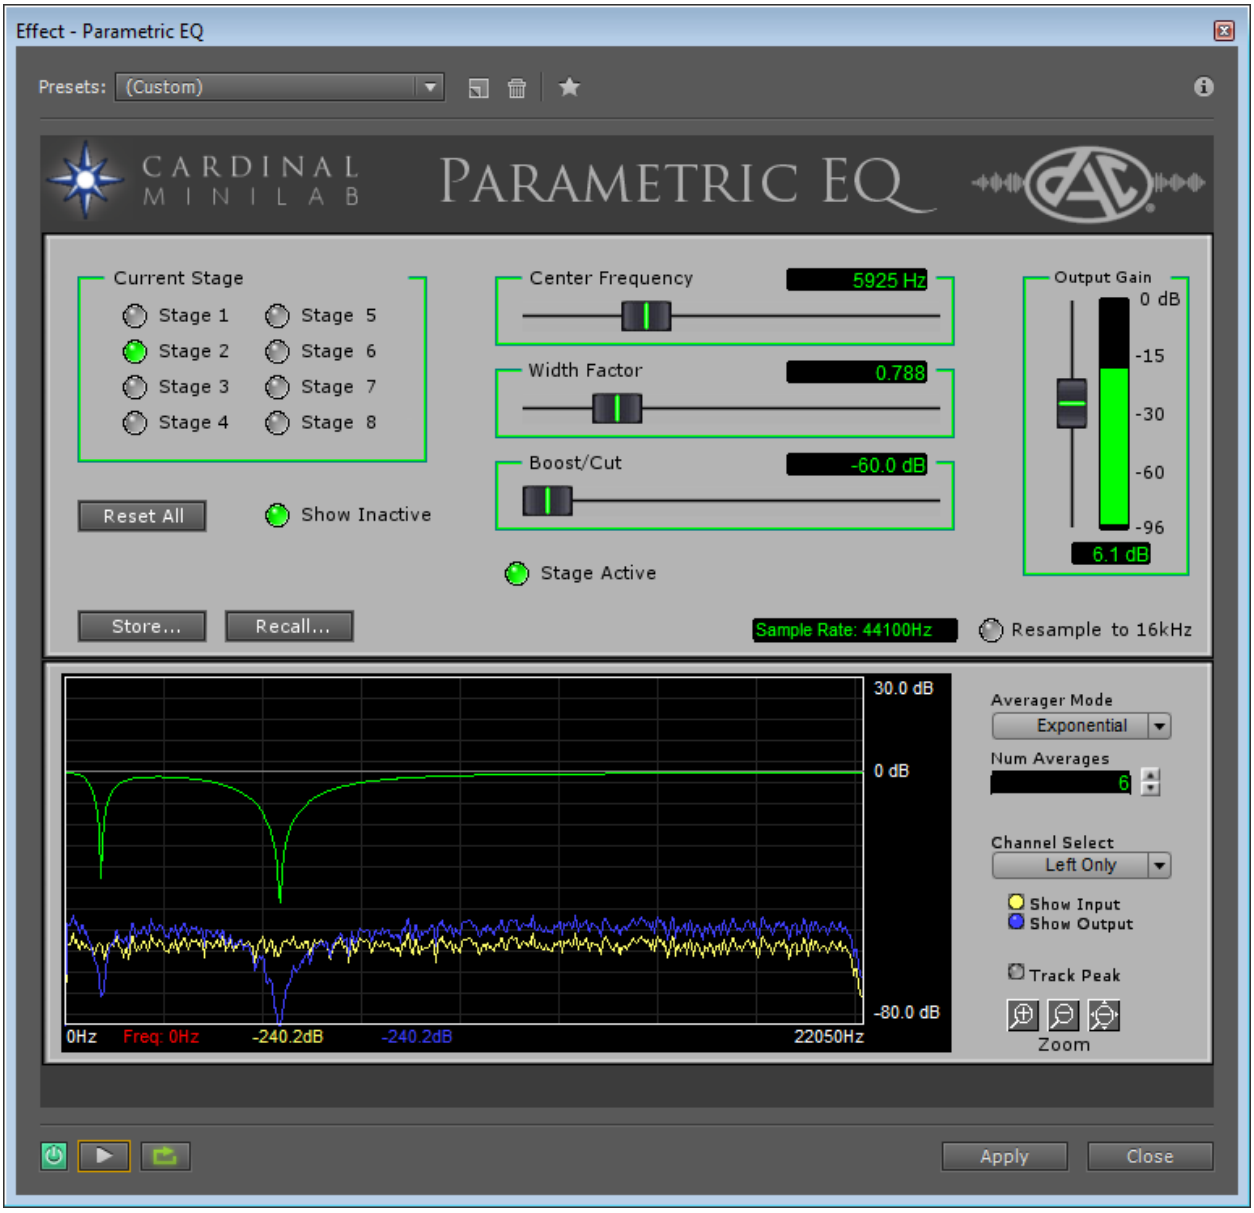

Figure 139: Parametric Equalizer Main Window

13.2 MAIN PLUG-IN WINDOW

Figure 139 above shows the plug-in as it would be displayed if used in Adobe Audition. The buttons above and below the plug-in window will appear differently depending upon your audio editing environment. Please consult the documentation on your specific audio editing environment for details of how to use these features.

This window provides access to all the functionality of the plug-in. A spectrum

analyzer is provided as an aid in determining the characteristics of the target audio. Details of the plug-in are described in the following sections.

13.3 FILTER CONTROLS

13.3.1 Current Stage

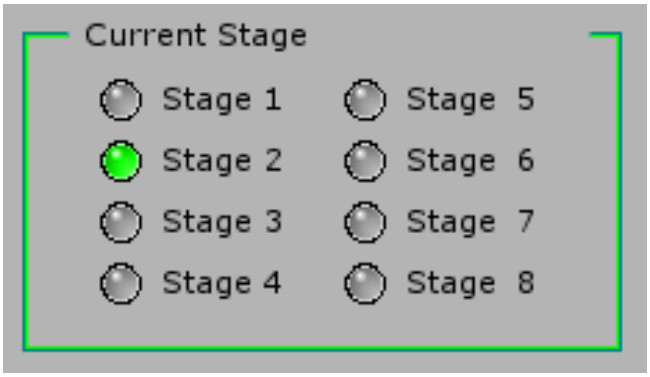

Figure 140: Current Stage

The buttons in the **Current Stage** box allow the user to select the current equalizer stage for adjustment. When a stage is selected, its settings populate the **Center Frequency, Width Factor, Boost/Cut** and **Stage Active** controls.

13.3.2 Center Frequency

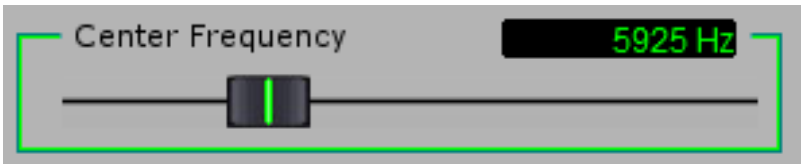

Figure 141: Center Frequency

The frequency at which the current stage's boost/cut region is centered.

13.3.3 Width Factor

Figure 142: Width Factor

A factor controlling the width of the current stage's boost/cut region.

13.3.4 Boost/Cut

Figure 143: Boost/Cut

The amount of boost or cut to be applied by the current stage.

13.3.5 Stage Active

Figure 144: Stage Active

If the indicator is lit, the current stage is being applied to audio. If the indicator is dark, the current stage is bypassed.

13.3.6 Output Gain

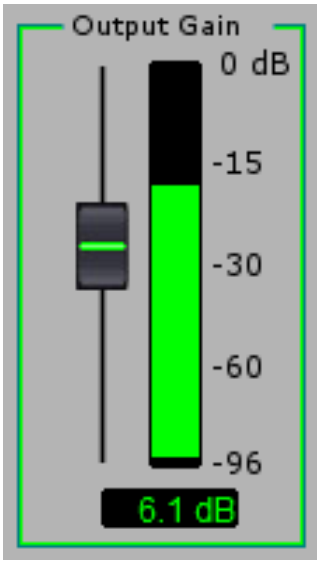

Figure 145: Output Gain

Amount of gain or attenuation applied to the audio after all active parametric EQ stages have been applied.

13.3.7 Reset All

Figure 146: Reset All

Returns all stages back to their default settings, which has no effect on the

audio (Center Frequency: 0Hz, Boost/Cut: 0dB, Stage Inactive).

13.3.8 Show Inactive

Figure 147: Show Inactive

Displays an orange trace on the graph to indicate what the transfer curve looks like *if* all the stages are enabled.

13.3.9 Store Button

Figure 148: Store Button

This button allows the user to store the current configuration to a user-specified disk file that will not be lost when the computer is turned off.

13.3.10 Recall Button

Figure 149: Recall Button

This button allows the user to recall a previously stored configuration from any of the saved disk files previously generated using the **Store** button.

14 FILTER CHAIN

14.1 OVERVIEW OF THE FILTER CHAIN

The Filter Chain plug-in provides the ability to put up to 8 MiniLab plug-ins in series and render the cumulative effect in one pass of the audio. Normally, each plug-in must be brought up separately in the editor and rendered on the audio independent from one another. This process can be time consuming and inefficient if you wish to apply more than one plug-in. With the Filter Chain, it is possible to create a chain of plug-ins, modify their settings and listen to their cumulative effect before rendering the audio just once.

Some audio editors provide their own “Effects Rack” or “Filter Chain”, but most do not. If you wish to chain together Cardinal MiniLab plug-ins, it is recommended that you use the Filter Chain plug-in to do it, as it is designed specifically to work with our plug-ins.

The spectrum analyzer displays the unmodified input from the audio editor and

the total resulting output audio after the last stage of the filter chain.

Figure 150: Filter Chain Main Window

14.2 MAIN PLUG-IN WINDOW

Figure 150 above shows the plug-in as it would be displayed if used in Adobe Audition. The buttons above and below the plug-in window will appear differently depending upon your audio editing environment. Please consult the documentation on your specific audio editing environment for details of how to use these features.

This window provides access to all the functionality of the plug-in. A spectrum analyzer is provided as an aid in determining the characteristics of the target audio. Details of the plug-in are described in the following sections.

14.3 FILTER CHAIN CONTROLS

14.3.1 Filter Select

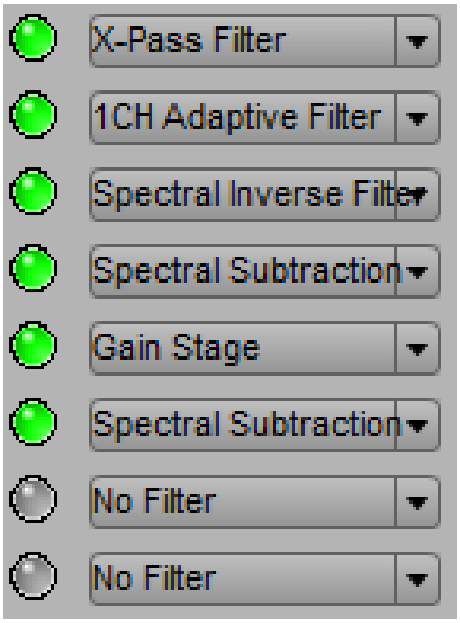

Figure 151: Filter Select Controls

The Filter Select controls are a set of 8 combo boxes that allow you to select from all the available Cardinal MiniLab plug-ins that are installed on your system. Some filters are only available if the audio presented matches their criteria (e.g. Reference Canceller is only available for stereo audio files).

The button next to each filter select control enables the filter. The audio always runs through each filter, however if the filter is disabled, its audio output is discarded. This allows adaptive filters to keep their solution in sync with the audio playback, but prevents the output of the filter to continue down the chain.

14.3.2 Filter Control Button Group

Figure 152: Filter Control Group

The Filter Control Group of buttons allows the user to move, delete and edit the corresponding plug-in.

The Move button will shift the plug-in up and down in the chain without altering its settings.

The Delete button will remove the corresponding plug-in from the chain and shift any filters below it up to fill the gap.

The Edit button brings up the graphical interface for the corresponding plug-in. You can have multiple plug-in windows open at one time.

14.3.3 Bypass All

Figure 153: Bypass All

The **Bypass All** control will disable all the filters in the chain. This is useful to easily hear before and after the filter chain processing to see if you are achieving the desired result.

14.3.4 Store Button

Figure 154: Store Button

The **Store** button will save the entire filter chain configuration, including all filter parameters, into a Filter Chain Settings (FCS) file for recall later. These files can also be used by the **Error! Reference source not found.** to apply the settings to a large number of audio files in batch mode.

14.3.5 Recall Button

Figure 155: Recall Button

The **Recall** button will restore the entire filter chain configuration from a Filter Chain Settings (FCS) file that was preserved using the **Store** button. Using this feature will wipe out the current Filter Chain configuration.

14.3.6 Clear Button

Figure 156: Clear Button

The **Clear** button will clear all the adaptive filters in the chain. It does not remove any filters from the chain.